
DiffSDA: A Framework for Unsupervised Diffusion Sequential Disentanglement Across Modalities

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Unsupervised representation learning, particularly sequential disentanglement,
2 aims to separate static and dynamic factors of variation in data without relying
3 on labels. This remains a challenging problem, as existing approaches based on
4 variational autoencoders and generative adversarial networks often rely on multiple
5 loss terms, complicating the optimization process. Furthermore, sequential disen-
6 tanglement methods face challenges when applied to real-world data, and there is
7 currently no established evaluation protocol for assessing their performance in such
8 settings. Recently, diffusion models have emerged as state-of-the-art generative
9 models, but no theoretical formalization exists for their application to sequential
10 disentanglement. In this work, we introduce the Diffusion Sequential Disentan-
11 glement Autoencoder (DiffSDA), a novel, modal-agnostic framework effective
12 across diverse real-world data modalities, including time series, video, and audio.
13 DiffSDA leverages a new probabilistic modeling, latent diffusion, and efficient
14 samplers, while incorporating a challenging evaluation protocol for rigorous test-
15 ing. Our experiments on diverse real-world benchmarks demonstrate that DiffSDA
16 outperforms recent state-of-the-art methods in sequential disentanglement.

17 1 Introduction

18 The advancements in diffusion models have demonstrated generative performance that surpasses
19 previous approaches, such as variational autoencoders (VAEs) and generative adversarial networks
20 (GANs) [1, 2, 3], earning recognition for their ability to produce high-quality samples across di-
21 verse data modalities. However, this remarkable performance often comes at the cost of requiring
22 large amounts of labeled data [4]. This reliance on labels underscores the growing importance of
23 unsupervised learning [5], which aims to unlock the potential of such models without the need for ex-
24 pensive annotations. Within unsupervised learning, *disentangled representation learning* has become
25 particularly significant [6]. This approach seeks to decompose latent representations into distinct
26 factors, where each factor captures a specific variation in the data. Such representations improve
27 interpretability [7], mitigate biases [8], and improve generalization [9]. A prominent challenge is
28 to develop a modal-agnostic approach for *sequential* data such as video, audio, and time series. In
29 particular, the goal is to decompose the sequential signal into separate static and dynamic latent com-
30 ponents in an unsupervised manner. For example, in a video of a person speaking, the static factors
31 could represent the person’s facial appearance, while the dynamic factors encode facial movements.
32 In audio recordings, static factors may correspond to the speaker’s identity, while dynamic factors
33 capture content of the speech.

34 Despite recent advancements, most sequential disentanglement methods [10, 11, 12, 13, 14, 15] rely
35 on VAEs and GANs, which often require complex optimization with extensive hyperparameter tuning.
36 For instance, C-DSVAE [12] requires *five* hyperparameters solely to balance its various loss terms.

Moreover, these models are often evaluated on toy datasets and struggle to produce high-quality samples in real-world scenarios. The reliance on VAEs and GANs is directly related to the absence of a modeling framework for sequential disentanglement within diffusion-based modeling. Further, existing diffusion architectures do not produce disentangled representations [16, 17]. We hypothesize that a diffusion-based framework can reduce hyperparameter tuning and improve sample quality, paving the way for more robust and scalable approaches to unsupervised sequential disentanglement.

In this work, we introduce *Diffusion Sequential Disentanglement Autoencoder (DiffSDA)*, a novel probabilistic framework for sequential disentanglement. Unlike prior tools [12, 14], our method models static and dynamic factors as interdependent, enhancing the expressivity of their marginal distributions. Notably, our approach is based on a **single** standard diffusion loss term, while producing high-quality results. Furthermore, DiffSDA is **modal-agnostic**, allowing it to disentangle data across diverse modalities, such as video, audio, and time series, with only minor adjustments to the network. This stands in contrast to modal-dependent methods, such as animation-based approaches for video, which rely on temporal and spatial consistency properties inherent to visual data [18], or methods designed specifically for audio that depend on spectral or temporal cues [19].

Practically, we implement a **sequential semantic encoder** and adopt the efficient sampling framework EDM [20]. Moreover, we incorporate a latent diffusion module (LDM) [3] into our architecture, which enables robust handling of high-dimensional, real-world data, outperforming prior sequential disentanglement methods. Finally, using our method, we demonstrate that applying principal component analysis (PCA) to the latent static and dynamic representations reveals a further disentanglement into multiple interpretable factors, showcasing the richness of the learned representations.

We perform a comprehensive evaluation of our model on standard benchmarks for sequential disentanglement [14] across three diverse data domains: audio, time series, and video. To further advance the field, we introduce a novel *evaluation protocol for high-quality visual sequential disentanglement*, incorporating three high-resolution video datasets and multiple quantitative metrics. Additionally, we propose a new post-training approach for disentangling representations into multiple factors. For the first time, our work presents a zero-shot task to demonstrate the generalizability of the factorization framework. Through these extensive evaluations, we show that DiffSDA not only effectively disentangles real-world data but also outperforms recent state-of-the-art methods. Our key contributions are summarized as follows:

1. We propose a novel modal-agnostic probabilistic framework for sequential disentanglement grounded in diffusion processes. Unlike most existing approaches, our formulation accommodates dependent static and dynamic factors of variation. The model is optimized using a single, unified score estimation loss.
2. Our design enables the effective disentanglement of high-dimensional, real-world data and supports zero-shot disentanglement tasks. Moreover, we demonstrate DiffSDA’s capability to disentangle static and dynamic information into multiple interpretable factors.
3. We provide a comprehensive evaluation demonstrating our model’s superiority in both qualitative and quantitative tasks, outperforming state-of-the-art methods. Additionally, we introduce a novel evaluation protocol specifically designed for video-based disentanglement.

2 Related Work

Generative modeling is a fundamental methodology for effectively sampling from numerical approximations of data distributions. Prominent approaches include variational autoencoders (VAEs) and generative adversarial networks (GANs) [21, 22]. More recently, diffusion models [23] and score matching [24, 25] have emerged as powerful alternatives, outperforming VAEs and GANs in generating high-quality samples through iterative denoising of latent variables [1, 2]. These methods are unified under a score-based modeling framework [26]. A critical challenge in generative modeling lies in representation learning, where semantic encodings of inputs are derived in an unsupervised manner. A related topic, center to this work, is the study of modal-agnostic disentangled representations, aiming to decompose data of various modalities into distinct factors of variation [6].

Disentangled Representation Learning. Most existing works on disentangled learning leverage VAEs and GANs to decompose non-sequential [27, 28, 29, 30, 31, 32] and sequential [33, 11, 34, 12,

13, 14, 15, 35, 36, 10] data. A key limitation of these approaches lies in their reliance on complex loss formulations, which typically involve multiple regularizers alongside the standard VAEs and GANs losses. While significant progress has been made in enhancing the generative capabilities of VAEs and GANs [37, 31], state-of-the-art methods for sequential disentanglement largely focus on simple datasets, far from real-world scenarios, with few exceptions like SPYL’s preliminary results [14]. In contrast, works in animation [18, 38, 39] have shown strong results on real-world data by leveraging video priors for disentangling objects and motion. However, these modal-dependent approaches can exploit relaxed assumptions and specialized tools, whereas our modal-agnostic method can adapt to diverse modalities, including video, audio, and time series.

Table 1: A comparison between animation, diffusion, and sequential disentanglement methods.

	Method	Modal Agnostic	Efficient	Real-World	Latent Factorization	Latents Prior	Loss Terms
ani- mation	FOM [18]	✗	✓	✓	✗	N/A	2
	AA [38]	✗	✓	✓	✗	N/A	1
	MA [39]	✗	✓	✓	✗	N/A	2
non seq.	DiffAE [16]	✗	✗	✓	✗	N/A	1
	InfoDiff [40]	✗	✗	✓	✗	N/A	2
sequen- tial	SPYL [14]	✓	✓	✗	✓	independent	5
	DBSE [15]	✓	✓	✗	✓	independent	2
	Ours	✓	✓	✓	✓	dependent	1

Diffusion-Based Disentanglement. The emergence of diffusion models has recently enabled novel approaches for non-sequential disentanglement [41, 42, 17, 43, 44, 45], achieving high-resolution image generation with disentangled factors. Moreover, other efforts have concentrated on structuring their latent representations. For instance, DiffAE [16] introduces an autoencoder to facilitate the manipulation of visual features, while InfoDiffusion [17] adds a loss regularizer to enhance disentanglement. Despite these advances, to the best of our knowledge, no theoretical formalization, and specifically, probabilistic modeling, has yet been proposed for diffusion-based disentanglement in sequential settings. Furthermore, practical approaches for this domain remain unexplored.

To contextualize our work within the landscape of existing tools, we present a comparative summary in Tab. 1, highlighting how our approach either advances or maintains all key aspects of representation learning. Specifically, while animation methods (FOM, AA, MA) and non-sequential diffusion tools (DiffAE, InfoDiff) handle real-world data, they are modal-dependent and do not provide a latent factorization. Within sequential disentanglement approaches (SPYL, DBSE), only our work supports real-world data via a single loss optimization.

3 Method

In this section, we introduce a novel probabilistic framework for unsupervised sequential disentanglement based on diffusion models. Currently, none of the existing approaches leverage diffusion models for unsupervised sequential disentanglement, leaving a significant gap in the field. Our framework addresses this gap by establishing a probabilistic modeling formalization and providing an efficient implementation for disentangling static and dynamic factors in sequential data. Background on diffusion models, diffusion autoencoders, and additional details about the method can be found in App. A and App. B. Throughout this section, and the subsequent ones, the subscripts represent time in the diffusion process, and superscripts indicate time in the sequence, e.g., a sequence state of the diffusion process is denoted by \mathbf{x}_t^τ , $t \in [0, T]$ and $\tau \in \{1, \dots, V\}$. T and V represent the maximum diffusion and sequence times, respectively. We consider discrete time sequences of continuous time diffusion processes; however, our modeling can be extended to additional settings.

3.1 Probabilistic Modeling

Existing frameworks for sequential disentanglement lack a probabilistic modeling foundation for diffusion-based modeling. To address this gap, we propose a novel probabilistic approach based on two diffusion models. The first model details the latent-independent distribution density of the static (time-invariant) and dynamic (time-variant) factors, \mathbf{s}_0 and $\mathbf{d}_0^{1:V}$, respectively. The second model

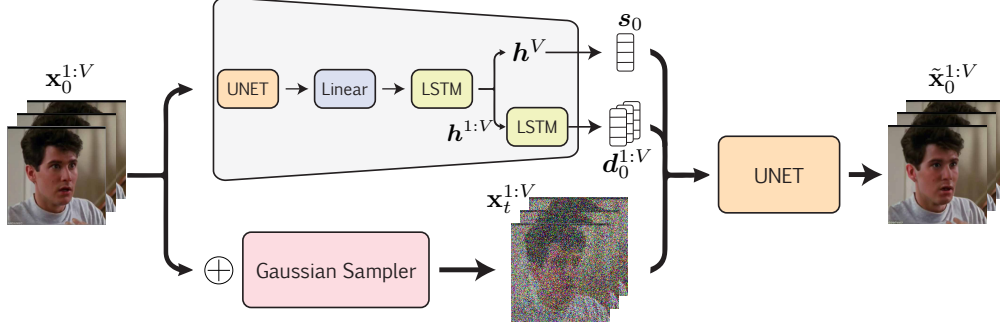


Figure 1: *DiffSDA* processes sequences $\mathbf{x}_0^{1:V}$ via semantic and stochastic encoders (top and bottom). Their outputs $(s_0, \mathbf{d}_0^{1:V}, \mathbf{x}_t^{1:V})$ are fed to a stochastic decoder yielding a denoised $\tilde{\mathbf{x}}_0^{1:V}$ (right).

129 specifies the observed distribution and its dependence on the disentangled factors. Formally, the joint
130 distribution is given by

$$p(\mathbf{x}_0^{1:V}, \mathbf{x}_T^{1:V}, s_0, s_T, \mathbf{d}_0^{1:V}, \mathbf{d}_T^{1:V}) = p_{T0}(s_0, \mathbf{d}_0^{1:V} \mid s_T, \mathbf{d}_T^{1:V}) \prod_{\tau=1}^V p_{T0}(\mathbf{x}_0^\tau \mid \mathbf{x}_T^\tau, s_0, \mathbf{d}_0^\tau) \quad (1)$$

131 where $p_{T0}(s_0, \mathbf{d}_0^{1:V} \mid s_T, \mathbf{d}_T^{1:V})$ is a standard diffusion process with $p_{T0}(\cdot)$ being the transition
132 distribution from time T to time 0 . The state distribution of $p_{T0}(\mathbf{x}_0^\tau \mid \mathbf{x}_T^\tau, s_0, \mathbf{d}_0^\tau)$ is conditioned on
133 the latent \mathbf{x}_T^τ and the factors s_0 and \mathbf{d}_0^τ .

134 Importantly, our probabilistic approach differs from existing work [12, 14] in that our static and
135 dynamic factors are interdependent. We motivate our model by three main reasons: i) expressiveness—
136 the overall dependence facilitates learning of different state trajectories, leading to higher expressivity
137 in the marginals $p_{t0}(\cdot)$; and ii) efficiency—our sampler is not autoregressive, allowing for fast and
138 parallelized sampling; and iii) causality—our model has the ability to learn intricate relationships
139 between the static and dynamic factors, if needed. We evaluate both the dependent and independent
140 approaches on our model to highlight the effectiveness of our approach. In summary, adopting
141 dependent modeling improves generation quality by 13%. Further details can be found in App. G.1.

142 Given a sequence $\mathbf{x}_0^{1:V} \sim p_0(\mathbf{x}_0^{1:V})$, the posterior distribution of the latent variables $\mathbf{x}_t^{1:V}$ and
143 latent factors s_0 and $\mathbf{d}_0^{1:V}$ is composed of three independent distributions. Further, unlike the
144 non-autoregressive prior in Eq. 1, here, we explicitly assume temporal dependence. The posterior
145 distribution reads

$$p(\mathbf{x}_t^{1:V}, s_0, \mathbf{d}_0^{1:V} \mid \mathbf{x}_0^{1:V}) = p_{0t}(\mathbf{x}_t^{1:V} \mid \mathbf{x}_0^{1:V}) p(s_0 \mid \mathbf{x}_0^{1:V}) \prod_{\tau=1}^V p(\mathbf{d}_0^\tau \mid \mathbf{d}_0^{<\tau}, \mathbf{x}_0^{<\tau}) \quad (2)$$

146 where $\mathbf{x}_t^{1:V}$ and s_0 are conditioned on the entire input $\mathbf{x}_0^{1:V}$, and the dynamic factors only depend on pr
147 evious dynamic factors and current and previous data elements. We employ score matching [24, 26],
148 to optimize for the denoising parametric map \mathbf{D}_θ . The map \mathbf{D}_θ takes the noisy latent \mathbf{x}_t^τ , time t , and
149 disentangled factors $\mathbf{z}_0^\tau := (s_0, \mathbf{d}_0^\tau)$, and it returns an estimate of the score function $\nabla_{\mathbf{x}} \log p_{0t}(\mathbf{x}_t^\tau \mid$
150 $\mathbf{x}_0^\tau)$. Overall, the optimization objective reads

$$\theta^* = \arg \min_{\theta} \mathbb{E}_t \left\{ \lambda_t \mathbb{E} \left[\|\mathbf{D}_\theta - \nabla_{\mathbf{x}} \log p_{0t}\|_2^2 \right] \right\}, \quad (3)$$

151 where $\lambda_t \in \mathbb{R}^+$ is a positive weight, $t \sim \mathcal{U}[0, T]$ is uniformly sampled over $[0, T]$, the variables
152 \mathbf{x}_t^τ , \mathbf{x}_0^τ are sampled from their respective distributions, $p_{0t}(\cdot)$, $p_0(\cdot)$, and \mathbf{z}_0^τ via the densities in Eq. 2.
153 The inner expectation is taken over \mathbf{x}_t^τ , \mathbf{z}_0^τ , and \mathbf{x}_0^τ . Importantly, p_{T0} of $s_0, \mathbf{d}_0^{1:V}$ is not used in Eq. 3,
154 and thus its optimization can be separated.

155 Notably, we make no assumptions about the given data $\mathbf{x}_0^{1:V}$, ensuring that our framework remains
156 modal-free and independent of specific properties of video, audio, or time series data. This theoretical
157 compatibility with any type of sequence makes it highly adaptable to diverse applications.

3.2 Diffusion Sequential Disentanglement Autoencoder

Our architecture, shown in Fig. 1, comprises three main components: (1) a sequential semantic encoder, (2) a stochastic encoder, and (3) a stochastic decoder. At a high level, the sequential semantic encoder factorizes data into separate static and dynamic components, while the stochastic decoder denoises the noisy latent representation produced by the stochastic encoder, conditioned on the disentangled factors. Notably, unlike prior works, our implementation achieves disentanglement with a single, simple loss term.

Encoders. Inspired by prior work in sequential disentanglement [11], we design a novel *sequential semantic encoder* to extract s_0 and $d_0^{1:V}$. Particularly, it consists of a U-Net [46] for video data and an MLP for other modalities, coupled with linear layers that independently process each sequence element. Then, an LSTM module summarizes the sequence into a latent representation $h^{1:V}$. The last hidden, h^V , is passed to a linear layer to produce s_0 , whereas $h^{1:V}$ are processed with another LSTM and a linear layer to produce $d_0^{1:V}$. Our stochastic encoder follows the EDM framework [20], adding noise $\epsilon \sim \mathcal{N}(0, \sigma_t^2 I)$ to each element \mathbf{x}_0^τ , yielding $\mathbf{x}_t^\tau = \mathbf{x}_0^\tau + \epsilon$. These encoders realize in practice the posterior in Eq. 2.

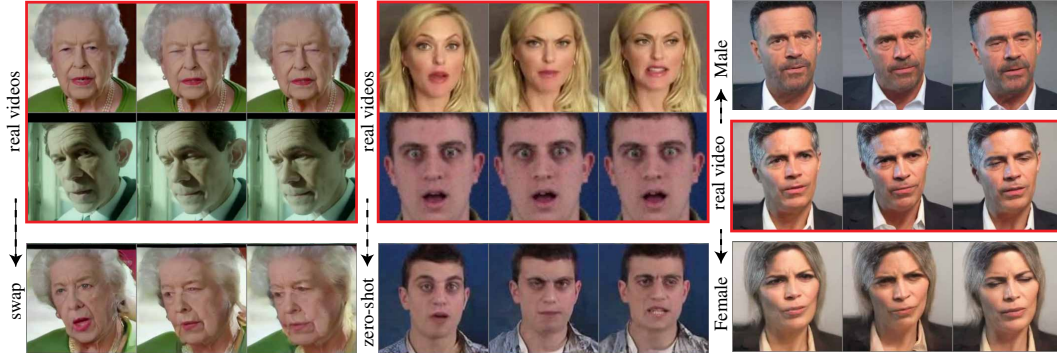


Figure 2: We present swap (left), zero-shot (middle), and multifactor disentanglement (right) results on multiple real-world and high-resolution visual datasets. See Sec. 4 for further details.

Decoder. To efficiently handle real-world sequential information, we follow the decoding in EDM [20], featuring only 63 neural function evaluations (NFEs) during inference. Our decoder D_θ takes as inputs the noisy input \mathbf{x}_t^τ and disentangled factors $\mathbf{z}_0^\tau := (s_0, d_0^\tau)$, and it returns a denoised version of \mathbf{x}_t^τ , denoted by $\tilde{\mathbf{x}}_0^\tau$. Given any $t \in [0, T]$ and $\tau \in \{1, \dots, V\}$, the decoder is parameterized independently from other times t', τ' as follows

$$\tilde{\mathbf{x}}_0^\tau := D_\theta(\mathbf{x}_t^\tau, t, \mathbf{z}_0^\tau) = c_t^{\text{skip}} \mathbf{x}_t^\tau + c_t^{\text{out}} \mathbf{F}_\theta(c_t^{\text{in}} \mathbf{x}_t^\tau, \mathbf{z}_0^\tau, c_t^{\text{noise}}), \quad (4)$$

where c_t^{skip} modulates the skip connection, $c_t^{\text{in}}, c_t^{\text{out}}$ scale the input/output magnitudes, and c_t^{noise} maps noise at time t into a conditioning input for the neural network \mathbf{F}_θ , conditioned on \mathbf{z}_0^τ through AdaGN.

Loss. While prior sequential disentanglement works depend on intricate prior modeling, regularization terms, and mutual information losses, leading to many hyper-parameters and challenging training, we opt for a simpler objective containing a single loss term that is based on Eq. 3,

$$\mathbb{E}_{t, \mathbf{x}_t^\tau, \mathbf{z}_0^\tau, \mathbf{x}_0^\tau} \left[\lambda_t (c_t^{\text{out}})^2 \|\mathbf{F}_\theta - \frac{1}{c_t^{\text{out}}} (\mathbf{x}_0^\tau - c_t^{\text{skip}} \cdot \mathbf{x}_t^\tau)\|_2^2 \right], \quad (5)$$

where \mathbf{F}_θ takes as inputs $c_t^{\text{in}} \mathbf{x}_t^\tau, \mathbf{z}_0^\tau$, and c_t^{noise} . While our loss in Eq. 5 does not include auxiliary terms, it promotes disentanglement due to two main reasons: i) the static factor s_0 is shared across τ , and thus it will not hold dynamic information, and ii) the dynamic factors $d_0^\tau \in \mathbb{R}^k$ are low-dimensional (i.e., k is small), making it difficult for d_0^τ to store static features. We empirically validate these assumptions through an experiment presented in App. G.2. Finally, we briefly mention that to support high-resolution sequences, we incorporate latent diffusion models (LDM) [3], using a pre-trained VQ-VAE autoencoder to reduce the high-dimensionality of input frames. Instead of factorizing all the equations above with new symbols for the features VQ-VAE produces, we denote by $\mathbf{x}_0^{1:V}$ the input sequence, and we abuse the notation $\mathbf{x}_0^{1:V}$ to denote the latent features, i.e., $\mathbf{x}_0^{1:V} = \mathcal{E}(\mathbf{x}_0^{1:V})$ and $\mathbf{x}_0^{1:V} = \mathcal{D}(\mathbf{x}_0^{1:V})$, where \mathcal{E} and \mathcal{D} are the VQ-VAE encoder and decoder, respectively.




Figure 3: We present dynamic swap results of our approach (third row) and SPYL (fourth row) on CelebV-HQ (left), VoxCeleb (middle), and TaiChi-HD (right).

4 Results

Below, we empirically evaluate the modeling capabilities of DiffSDA in comparison to recent *modal-agnostic* state-of-the-art methods (see Tab. 1), SPYL [14] and DBSE [15]. In general, we consider quantitative and qualitative experiments. For video, we include three high-resolution, real-world visual datasets that have not been previously used for sequential disentanglement: VoxCeleb [47], CelebV-HQ [48], and TaiChi-HD [18], along with the popular MUG dataset [49]. For audio, we consider TIMIT [50] and a new dataset, Libri Speech [51]. The time series datasets are PhysioNet, ETTh1, and Air Quality [52]. Detailed descriptions of the datasets and their pre-processing can be found in App. D. For brevity, we omit below the subscript indicating the diffusion step for clean samples (corresponding to time step 0).

4.1 Conditional swap in videos

We begin our tests with the conditional swap task [11]. Given two sample videos $\mathbf{x}, \hat{\mathbf{x}} \sim p_0$, the goal in this experiment is to create a new sample $\bar{\mathbf{x}}$, conditioned on the static factor of \mathbf{x} and dynamic features of $\hat{\mathbf{x}}$. This is done by extracting the latent factors $\mathbf{z} = (s, \mathbf{d}^{1:V})$ and $\hat{\mathbf{z}} = (\hat{s}, \hat{\mathbf{d}}^{1:V})$ for \mathbf{x} and $\hat{\mathbf{x}}$, respectively. The new sample $\bar{\mathbf{x}}$ is defined to be the reconstruction of $\bar{\mathbf{z}} = (s, \hat{\mathbf{d}}^{1:V})$ through sampling, see Alg. 1. In an ideal swap, $\bar{\mathbf{x}}$ preserves the static characteristics of \mathbf{x} while presenting the dynamics of $\hat{\mathbf{x}}$, thus demonstrating strong disentanglement capabilities of the swapping method. We show in Fig. 2 (left) a swap example of DiffSDA, where the top two rows are real videos, and the third row shows the new sample obtained by preserving the static features of the first row and using the dynamics of the second row. Remarkably, while the people in these sequences are very different, many fine details are transferred, including head angle and orientation, as well as mouth and eyes orientation and openness. In Fig. 3, we present additional swap results on CelebV-HQ (left), VoxCeleb (middle), and TaiChi-HD (right), comparing DiffSDA (third row) to SPYL (fourth row). Notably, our approach produces high-quality samples, while swapping the dynamics of the second row into the first row, whereas SPYL struggles both with the reconstruction and swap. Additional conditional and unconditional swap results appear in App. H.2 and App. H.3, respectively.

In addition to the above qualitative evaluation, we also want to quantitatively assess DiffSDA’s effectiveness. We report in App. F results from the traditional quantitative benchmark, where a pre-trained judge (classifier) is used to determine if swapped content is correct [12]. However, there are two main issues with the benchmark: i) it depends on labeled data, making it relevant to only a small number of datasets; and ii) results are sensitive to the expressivity and generalizability of the judge. For instance, swapping a smiling expression from person A to person B, may result in person B having a smile, different from the one in the data. In these cases, the judge may wrongly classify a different expression to the smiling person B, see App. F for further discussion.

Towards addressing these issues, we propose new *unsupervised* swapping metrics to quantitatively measure the model’s disentanglement abilities. We adopt estimators commonly used in animation for assessing whether objects and motions are preserved [18]. Specifically, we utilize the *average Euclidean distance* (AED) that is based on the distances between the latent representations of




Figure 4: Zero-shot swap results, training on VoxCeleb and tested on CelebV-HQ or MUG.

images. Further, we also employ the *average keypoint distance* (AKD) which computes the distances between selected keypoints in images. Intuitively, AED and AKD have been designed to identify the preservation of objects and motions in images, respectively. See App. E for definitions.

Equipped with these new metrics, we perform conditional swapping over a pre-defined random list of sample pairs, $\mathbf{x}, \hat{\mathbf{x}}$. Particularly, we reconstruct new samples of the form $\mathbf{z}^s := (s, \hat{\mathbf{d}}^{1:V})$ and $\mathbf{z}^d := (\hat{s}, \mathbf{d}^{1:V})$, encoding dynamic and static swaps, respectively. We compute the AED of \mathbf{z}^s with respect to \mathbf{z} (arising from \mathbf{x}), expecting their static features to be similar. Following the same logic, we compute the AKD of \mathbf{x}^d (reconstructed from \mathbf{z}^d) and \mathbf{x} , as they share the dynamic factors. Our findings are presented in Tab. 2, where DiffSDA outperforms SOTA previous (SPYL, DBSE) approaches across all datasets, except for AED on TaiChi-HD, where we attain the second best error. Notably, our AKD errors are significantly lower than SPYL and DBSE. Further, we apply these metrics to assess reconstruction performance, as well as the mean squared error (MSE), with the results shown in Tab. 3. Again, DiffSDA is superior to current SOTA methods. Additionally, we include a generative evaluation in App. G.4, comparing our approach to previous methods.

4.2 Zero-shot video disentanglement

In the previous sub-section, the conditional swap was performed on the held-out test set of each dataset on which we trained on. In contrast to previous work, for the first time, we perform the same task on a dataset unseen during training. We show an example in Fig. 2 (middle) of zero-shot swap, where our model was trained on the VoxCeleb dataset (1st row) and the inferred sequence was taken from MUG (2nd row). Particularly, we froze the static features of the MUG sample and swapped the dynamic factors with those of VoxCeleb (3rd row). Remarkably, in addition to changing the facial expression of the person, DiffSDA also adds the necessary details to mimic the body pose. We emphasize that the MUG dataset does not include sequences similar to the third row in Fig. 2, but rather zoomed-in facial videos as shown in the second row, thus, our zero-shot results present a significant adaptation to the new data. Additionally, we include in Fig. 4 zero-shot examples where DiffSDA is trained on VoxCeleb and evaluated on CelebV-HQ or MUG. These results further highlight the effectivity of our approach in transferring dynamic features across different datasets. Finally, we provide more zero-shot examples in App. H.4.

Table 2: Preservation of objects (AED) and motions (AKD) is estimated across several datasets and methods. The labels ‘static frozen’ and ‘dynamics frozen’ correspond to samples \mathbf{z}^s and \mathbf{z}^d .

	AED↓ (static frozen)			AKD↓ (dynamics frozen)		
	SPYL	DBSE	Ours	SPYL	DBSE	Ours
MUG (64×64)	0.766	0.773	0.751	1.132	1.118	0.802
VoxCeleb (256×256)	1.058	1.026	0.846	4.705	10.96	2.793
CelebV-HQ (256×256)	0.631	0.751	0.540	39.16	28.69	6.932
TaiChi-HD (64×64)	0.443	0.325	0.326	7.681	6.312	2.143

4.3 Toward multifactor video disentanglement

Multifactor sequential disentanglement is a challenging problem, where the objective is to produce several static factors and several dynamic factors per frame [53]. Here, we show that our model has the potential to further disentangle the static and dynamic features into additional factors of variation. Inspired by DiffAE [16], we explore the learned latent space in an unsupervised linear fashion, particularly, using principal component analysis (PCA). Namely, to obtain fine-grained semantic static factors of variation, we sample a large batch of static vectors $\hat{s}_j \in \mathbb{R}^h$, with h the static latent size, $j = 1, \dots, b = 2^{15}$. Then, we compute PCA on the matrix formed by arranging $\{\hat{s}_j\}$ in its columns, yielding the principal components $\{v_i\}_{i=1}^h$, given that $b \geq h$. We can utilize the latter pool of static variability by exploring the latent space from a static code s of a real example \mathbf{x} in the test set, i.e.,

$$\bar{s} = \left(\frac{s - \mu_{\hat{s}}}{\sigma_{\hat{s}}} + \alpha v_i \cdot \sqrt{h} \right) \cdot \sigma_{\hat{s}} + \mu_{\hat{s}}, \quad (6)$$

where $\mu_{\hat{s}}$ and $\sigma_{\hat{s}}^2$ are the mean and variance of the sampled static features, $\{\hat{s}_j\}_{j=1}^b$, and $\alpha \in [-\kappa, \kappa]$, notice that $\alpha = 0$ recovers the original sequence. The new sample $\bar{\mathbf{x}}$ is obtained by reconstructing the new static features \bar{s} with the original dynamic factors $\mathbf{d}^{1:V}$ of \mathbf{x} .

We demonstrate a static PCA exploration in Fig. 2 (right) on VoxCeleb. The middle row is the real video, whereas the top and bottom rows use positive and negative α values, respectively. Our results show that traversing in the positive direction yields more masculine appearances, and in contrast, going in the negative direction produces more feminine characters. Importantly, we highlight that other static features and the dynamics are fully preserved across the sequence. In App. H.5, we present further results on full sequences using multiple α values to demonstrate the gradual transition in the latent space. Notably, we find in our exploration principal components that control other features such as skin tone, image blurriness, and more.

4.4 Speaker identification in audio

Our approach is inherently modal-agnostic and extends beyond the video domain. Unlike methods tailored specifically for video or audio, which often require extensive modifications when applied to new modalities, our method is versatile and can adapt to different modalities with minimal adjustments to the backbone architecture. For example, to process audio data, we simply replace the U-Net with an MLP. In Tab. 4, we demonstrate the adaptability of our model by successfully disentangling audio data from the TIMIT dataset and Libri Speech, where TIMIT is a widely used benchmark for speech-related tasks and Libri Speech is an additional dataset we add for this benchmark. Following the speaker identification benchmarks [11], we evaluate disentanglement quality using the Equal Error Rate (EER), a standard metric in speech tasks. Specifically, the Static EER measures how effectively the static latent representations capture speaker identity, and similarly, the Dynamic EER assesses the dynamic latent representations. Notably, a well-disentangled model should yield a low Static EER (capturing speaker identity in static representations) and a high Dynamic EER (capturing content-related dynamics without speaker identity). The overall goal is to maximize the gap between these two metrics (Dis. Gap). Our model, achieves in TIMIT a disentanglement gap improvement of over 11%, with a 42.29% compared to 31.11% achieved by DBSE, thereby surpassing current state-of-the-art methods. Similar strong performance is achieved on Libri Speech as well. These results highlight the efficacy of our approach in the audio domain. Additional details regarding the dataset, evaluation metrics, and implementation are provided in the appendix. Furthermore, we report speech quality and reconstruction results in App. G.3, further validating our model’s effectiveness in the audio domain.

Table 3: Reconstruction errors are measured in terms of AED, AKD, and MSE across several datasets and models. We find DiffSDA to be orders-of-magnitude better than other methods.

	AED↓			AKD↓			MSE↓		
	SPYL	DBSE	Ours	SPYL	DBSE	Ours	SPYL	DBSE	Ours
MUG	0.49	0.49	0.11	0.47	0.48	0.06	0.001	0.001	3e−7
VoxCeleb	0.99	1.03	0.37	2.27	2.43	1.09	0.005	0.003	5e−4
CelebV-HQ	0.70	0.78	0.29	15.0	13.8	1.26	0.012	0.006	6e−4
TaiChi-HD	0.32	0.29	0.001	4.31	3.83	0.10	0.018	0.007	2e−7

Table 4: Disentanglement metrics on TIMIT and LibriSpeech

Method	TIMIT			LibriSpeech		
	Static EER↓	Dynamic EER↑	Dis. Gap↑	Static EER↓	Dynamic EER↑	Dis. Gap↑
DSVAE	5.64%	19.20%	13.56%	15.06%	28.94%	13.87%
SPYL	3.41%	33.22%	29.81%	24.87%	49.76%	24.89%
DBSE	3.50%	34.62%	31.11%	16.75%	22.61%	5.58%
Ours	4.43%	46.72%	42.29%	11.02%	45.94%	34.93%

4.5 Downstream prediction and classification tasks on time series information

Finally, we evaluate our approach on time series data, following the evaluation protocol in [15]. The evaluation is carried out in two main independent setups: 1) We assess the quality of the learned latent representations using a predictive task. The model is trained on a dataset, and at test time, the static and dynamic factors are extracted and used as input features for a predictive model. Two tasks are considered: (i) predicting mortality risk with the PhysioNet dataset [54], and (ii) predicting oil temperature using the ETTh1 dataset [55]. Performance is evaluated using AUPRC and AUROC for PhysioNet, and Mean Absolute Error (MAE) for ETTh1. 2) We investigate the model’s ability to capture global patterns within its disentangled static latent representations, which have been shown to enhance performance [56]. Following a similar procedure, the model is trained, and now only the static representations are extracted. These representations are then used as input features for a classifier. For the PhysioNet dataset, Intensive Care Unit (ICU) unit types are used as global labels, while for the Air Quality dataset, the month of the year serves as the target variable. Further details regarding datasets, metrics, and implementation can be found in App. D and App. E. We compare our method vs. state-of-the-art baselines, including DBSE, SPYL, and GLR [52]. Results for predictive and classification tasks are given in Tab. 5. Notably, our model outperforms across all tasks.

Table 5: Time series prediction and classification benchmarks.

Task		GLR	SPYL	DBSE	Supervised	Ours
pred.	AUPRC↑	0.37 ± 0.09	0.37 ± 0.02	0.47 ± 0.02	0.44 ± 0.02	0.50 ± 0.006
	AUROC↑	0.75 ± 0.01	0.76 ± 0.04	0.86 ± 0.01	0.80 ± 0.04	0.87 ± 0.004
	MAE↓ (ETTh1)	12.3 ± 0.03	12.2 ± 0.03	11.2 ± 0.01	10.19 ± 0.20	9.89 ± 0.280
cls.	PhysioNet↑	38.9 ± 2.48	47.0 ± 3.04	56.9 ± 0.34	62.00 ± 2.10	64.6 ± 0.35
	Air Quality↑	50.3 ± 3.87	57.9 ± 3.53	65.9 ± 0.01	62.43 ± 0.54	69.2 ± 1.50

5 Conclusions

The analysis and results of this study underscore the potential of the proposed DiffSDA model to address key limitations in sequential disentanglement, specifically in the context of complex real-world visual data, speech audio, and time series. By leveraging a novel probabilistic framework, diffusion autoencoders, efficient samplers, and latent diffusion models, DiffSDA provides a robust solution for disentangling both static and dynamic factors in sequences, outperforming existing state-of-the-art methods. Moreover, the introduction of a new real-world visual evaluation protocol marks a significant step towards standardizing the assessment of sequential disentanglement models. Nevertheless, while DiffSDA shows promise in handling high-resolution videos and varied datasets, future research should focus on optimizing its computational efficiency and extending its applicability to more diverse sequence modalities, such as sensor data. Such modalities present unique challenges, as varying temporal characteristics and distinct data patterns, which may require adapting the model architecture and training strategies. Finally, a key challenge ahead lies in fully extending our multifactor exploration procedure to effectively disentangle and represent multiple interacting factors [53]. We leave these considerations and further explorations for future work.

References

- [1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [2] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [3] 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*, pages 10684–10695, 2022.
- [4] Tianhong Li, Dina Katabi, and Kaiming He. Return of unconditional generation: A self-supervised representation generation method. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.
- [5] Yoshua Bengio, Aaron C Courville, and Pascal Vincent. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR, abs/1206.5538*, 1(2665):2012, 2012.
- [6] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [7] Wenqian Liu, Runze Li, Meng Zheng, Srikrishna Karanam, Ziyang Wu, Bir Bhanu, Richard J Radke, and Octavia I Camps. Towards visually explaining variational autoencoders. in 2020 ieee. In *CVF Conference on Computer Vision and Pattern Recognition, CVPR*, pages 13–19, 2020.
- [8] Elliot Creager, David Madras, Jörn-Henrik Jacobsen, Marissa Weis, Kevin Swersky, Toniann Pitassi, and Richard Zemel. Flexibly fair representation learning by disentanglement. In *International conference on machine learning*, pages 1436–1445. PMLR, 2019.
- [9] Hanlin Zhang, Yi-Fan Zhang, Weiyang Liu, Adrian Weller, Bernhard Schölkopf, and Eric P Xing. Towards principled disentanglement for domain generalization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8024–8034, 2022.
- [10] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018.
- [11] Li Yingzhen and Stephan Mandt. Disentangled sequential autoencoder. In Jeifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 5670–5679. PMLR, 10–15 Jul 2018.
- [12] Junwen Bai, Weiran Wang, and Carla P Gomes. Contrastively disentangled sequential variational autoencoder. *Advances in Neural Information Processing Systems*, 34:10105–10118, 2021.
- [13] Jun Han, Martin Renqiang Min, Ligong Han, Li Erran Li, and Xuan Zhang. Disentangled recurrent wasserstein autoencoder. In *9th International Conference on Learning Representations, ICLR*, 2021.
- [14] Ilan Naiman, Nimrod Berman, and Omri Azencot. Sample and predict your latent: Modality-free sequential disentanglement via contrastive estimation. In *International Conference on Machine Learning*, pages 25694–25717. PMLR, 2023.
- [15] Nimrod Berman, Ilan Naiman, Idan Arbiy, Gal Fadlon, and Omri Azencot. Sequential disentanglement by extracting static information from a single sequence element. In *Forty-first International Conference on Machine Learning*, 2024.
- [16] Konpat Preechakul, Nattanat Chatthee, Suttisak Wizadwongsa, and Supasorn Suwajanakorn. Diffusion autoencoders: Toward a meaningful and decodable representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10619–10629, 2022.

- [17] Yingheng Wang, Yair Schiff, Aaron Gokaslan, Weishen Pan, Fei Wang, Christopher De Sa, and Volodymyr Kuleshov. InfoDiffusion: Representation learning using information maximizing diffusion models. In *International Conference on Machine Learning*, pages 36336–36354. PMLR, 2023.
- [18] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. First order motion model for image animation. In *Conference on Neural Information Processing Systems (NeurIPS)*, December 2019.
- [19] Sicheng Xu, Guojun Chen, Yu-Xiao Guo, Jiaolong Yang, Chong Li, Zhenyu Zang, Yizhong Zhang, Xin Tong, and Baining Guo. Vasa-1: Lifelike audio-driven talking faces generated in real time. *arXiv preprint arXiv:2404.10667*, 2024.
- [20] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.
- [21] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [22] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [23] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.
- [24] Aapo Hyvärinen and Peter Dayan. Estimation of non-normalized statistical models by score matching. *Journal of Machine Learning Research*, 6(4), 2005.
- [25] Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computation*, 23(7):1661–1674, 2011.
- [26] 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*, 2021.
- [27] Irina Higgins, Loic Matthey, Arka Pal, Christopher P Burgess, Xavier Glorot, Matthew M Botvinick, Shakir Mohamed, and Alexander Lerchner. beta-VAE: Learning basic visual concepts with a constrained variational framework. *ICLR (Poster)*, 3, 2017.
- [28] Ricky TQ Chen, Xuechen Li, Roger B Grosse, and David K Duvenaud. Isolating sources of disentanglement in variational autoencoders. *Advances in neural information processing systems*, 31, 2018.
- [29] Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In *International conference on machine learning*, pages 2649–2658. PMLR, 2018.
- [30] Luan Tran, Xi Yin, and Xiaoming Liu. Disentangled representation learning GAN for pose-invariant face recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1415–1424, 2017.
- [31] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of styleGAN. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8110–8119, 2020.
- [32] Xuanchi Ren, Tao Yang, Yuwang Wang, and Wenjun Zeng. Learning disentangled representation by exploiting pretrained generative models: A contrastive learning view. *arXiv preprint arXiv:2102.10543*, 2021.
- [33] Wei-Ning Hsu, Yu Zhang, and James Glass. Unsupervised learning of disentangled and interpretable representations from sequential data. *Advances in neural information processing systems*, 30, 2017.

- [34] Yizhe Zhu, Martin Renqiang Min, Asim Kadav, and Hans Peter Graf. S3VAE: self-supervised sequential VAE for representation disentanglement and data generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6538–6547, 2020.
- [35] Mathieu Cyrille Simon, Pascal Frossard, and Christophe De Vleeschouwer. Sequential representation learning via static-dynamic conditional disentanglement. In *European Conference on Computer Vision*, pages 110–126. Springer, 2025.
- [36] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *arXiv preprint arXiv:1706.08033*, 2017.
- [37] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. *Advances in neural information processing systems*, 33:19667–19679, 2020.
- [38] Li Hu. Animate anyone: Consistent and controllable image-to-video synthesis for character animation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8153–8163, 2024.
- [39] Zhongcong Xu, Jianfeng Zhang, Jun Hao Liew, Hanshu Yan, Jia-Wei Liu, Chenxu Zhang, Jiashi Feng, and Mike Zheng Shou. Magicanimate: Temporally consistent human image animation using diffusion model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1481–1490, 2024.
- [40] Jun Wang, Yinglu Liu, Yibo Hu, Hailin Shi, and Tao Mei. Facex-zoo: A pytorch toolbox for face recognition. In *Proceedings of the 29th ACM international conference on Multimedia*, pages 3779–3782, 2021.
- [41] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. *arXiv preprint arXiv:2210.10960*, 2022.
- [42] Tao Yang, Yuwang Wang, Yan Lu, and Nanning Zheng. DisDiff: Unsupervised disentanglement of diffusion probabilistic models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [43] Tao Yang, Cuiling Lan, Yan Lu, et al. Diffusion model with cross attention as an inductive bias for disentanglement. *arXiv preprint arXiv:2402.09712*, 2024.
- [44] Ye Zhu, Yu Wu, Zhiwei Deng, Olga Russakovsky, and Yan Yan. Boundary guided learning-free semantic control with diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024.
- [45] Stefan Andreas Baumann, Felix Krause, Michael Neumayr, Nick Stracke, Vincent Tao Hu, and Björn Ommer. Continuous, subject-specific attribute control in T2I models by identifying semantic directions. *arXiv preprint arXiv:2403.17064*, 2024.
- [46] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pages 234–241. Springer, 2015.
- [47] Arsha Nagrani, Joon Son Chung, and Andrew Senior. VoxCeleb: A large-scale speaker identification dataset. In *18th Annual Conference of the International Speech Communication Association, Interspeech*, pages 2616–2620. ISCA, 2017.
- [48] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. CelebV-HQ: A large-scale video facial attributes dataset. In *ECCV*, 2022.
- [49] Niki Aifanti, Christos Papachristou, and Anastasios Delopoulos. The MUG facial expression database. In *11th International Workshop on Image Analysis for Multimedia Interactive Services WIAMIS 10*, pages 1–4, 2010.
- [50] John S Garofolo. TIMIT acoustic phonetic continuous speech corpus. *Linguistic Data Consortium*, 1993, 1993.

- [51] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: an asr corpus based on public domain audio books. In *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pages 5206–5210. IEEE, 2015.
- [52] Sana Tonekaboni, Chun-Liang Li, Serkan O Arik, Anna Goldenberg, and Tomas Pfister. Decoupling local and global representations of time series. In *International Conference on Artificial Intelligence and Statistics*, pages 8700–8714. PMLR, 2022.
- [53] Nimrod Berman, Ilan Naiman, and Omri Azencot. Multifactor sequential disentanglement via structured Koopman autoencoders. In *The Eleventh International Conference on Learning Representations*, 2023.
- [54] Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220, 2000.
- [55] Shuyi Zhang, Bin Guo, Anlan Dong, Jing He, Ziping Xu, and Song Xi Chen. Cautionary tales on air-quality improvement in Beijing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 473(2205):20170457, 2017.
- [56] Shubhendu Trivedi, Zachary A Pardos, and Neil T Heffernan. The utility of clustering in prediction tasks. *arXiv preprint arXiv:1509.06163*, 2015.
- [57] Brian DO Anderson. Reverse-time diffusion equation models. *Stochastic Processes and their Applications*, 12(3):313–326, 1982.
- [58] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018.
- [59] Alexey Dosovitskiy and Thomas Brox. Generating images with perceptual similarity metrics based on deep networks. In D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 29, 2016.
- [60] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [61] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [62] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR, 2014.
- [63] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [64] Ali Razavi, Aaron Van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. *Advances in neural information processing systems*, 32, 2019.
- [65] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. Analyzing and improving the training dynamics of diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24174–24184, 2024.
- [66] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [67] Samuel Albanie, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Emotion recognition in speech using cross-modal transfer in the wild. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 292–301, 2018.

- 523 [68] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face
524 alignment problem? (and a dataset of 230,000 3d facial landmarks). In *Proceedings of the IEEE*
525 *International Conference on Computer Vision (ICCV)*, Oct 2017.
- 526 [69] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose
527 estimation using part affinity fields. In *Proceedings of the IEEE Conference on Computer Vision*
528 *and Pattern Recognition (CVPR)*, July 2017.
- 529 [70] Sefik Ilkin Serengil and Alper Ozpinar. Lightface: A hybrid deep face recognition framework.
530 In *2020 Iovations in Intelligent Systems and Applications Conference (ASYU)*, pages 23–27.
531 IEEE, 2020.
- 532 [71] Alexander Hermans, Lucas Beyer, and Bastian Leibe. In defense of the triplet loss for person
533 re-identification. *arXiv preprint arXiv:1703.07737*, 2017.
- 534 [72] Chandan KA Reddy, Vishak Gopal, and Ross Cutler. Dnsmos: A non-intrusive perceptual
535 objective speech quality metric. In *ICASSP*, 2021.
- 536 [73] Andreas Blattmann, Robin Rombach, Huan Ling, Tim Dockhorn, Seung Wook Kim, Sanja
537 Fidler, and Karsten Kreis. Align your latents: High-resolution video synthesis with latent
538 diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*
539 *Recognition*, pages 22563–22575, 2023.

A Background

A.1 Diffusion Models

Diffusion models [23] are a family of SOTA generative models, that were recently described using stochastic differential equations (SDEs), diffusion processes, and score-based modeling [26]. We will use diffusion models and score-based models interchangeably. These models include two processes: the forward process and the reverse process. The forward process (often not learnable) is an iterative procedure that corrupts the data by progressively adding noise to it. Specifically, the change to the state \mathbf{x}_t can be formally described by

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t)dt + g(t)d\mathbf{w}, \quad (7)$$

where \mathbf{w} is the standard Wiener process, $\mathbf{f}(\cdot, t)$ is a vector-valued function called the drift coefficient, and $g(\cdot)$ is a scalar function known as the diffusion coefficient. From a probabilistic viewpoint, Eq. 7 is associated with modeling the transition from the given data distribution, $\mathbf{x}_0 \sim p_0$, to p_t , the probability density of \mathbf{x}_t , $t \in [0, T]$. Typically, the prior distribution p_T is a simple Gaussian distribution with fixed mean and variance that contains no information of p_0 . The reverse process, which is learnable, de-noises the data iteratively. The reverse of a diffusion process is also a diffusion process, depending on the score function $\nabla_{\mathbf{x}} \log p_t(\mathbf{x})$ and operating in reverse time [57]. In our approach, we utilize the conditioned reverse process

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - g(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{u})]d\bar{t} + g(t)d\bar{\mathbf{w}}, \quad (8)$$

where $\bar{\mathbf{w}}$ is a standard Wiener process as time progresses backward from T to 0, $d\bar{t}$ is an negative timestep, and \mathbf{u} is a condition variable. Diffusion models are generative by sampling from p_T and use $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{u})$ to iteratively solve Eq. 8 until samples from p_0 are recovered.

A.2 Diffusion Autoencoders

Although diffusion models are powerful generative tools, they are not inherently designed to learn meaningful representations of the data. To address this limitation, several works [16, 17] have adapted diffusion models into autoencoders, resulting in diffusion autoencoders (DiffAEs). These models have demonstrated the ability to learn semantic representations of the data, allowing certain modifications of the resulting samples by altering their latent vectors. To this end, DiffAEs introduce a semantic encoder, taking a data sample x_0 and returning its semantic latent encoding z_{sem} . Then, the latter vector conditions the reverse process, enhancing the model’s ability to reconstruct and manipulate data samples. In practice, the denoiser is also conditioned on a feature map h and the time t , combined using an adaptive group normalization (AdaGN) layer [2]. The AdaGN block is defined as

$$\text{AdaGN}(h, t, z_{\text{sem}}) = z_s (t_s \text{GroupNorm}(h) + t_b), \quad (9)$$

where z_s is the output of a linear layer applied to z_{sem} , t_s and t_b are the outputs of a multi-layer perceptron (MLP) applied to the time t , and multiplications are done element-wise.

B DiffSDA Modeling

B.1 Unsupervised Sequential Disentanglement

Unsupervised sequential disentanglement is a challenging problem in representation learning, aiming to decompose a given dataset to its static (time-independent) and dynamic (time-dependent) factors of variation. Let $\mathcal{D} = \{\mathbf{x}_j^{1:V}\}_{j=1}^N$ be a dataset with N sequences $\mathbf{x}_j^{1:V} := \{\mathbf{x}_j^1, \dots, \mathbf{x}_j^V\}$, where $\mathbf{x}_j^T \in \mathbb{R}^d$. We omit the subscript j for brevity, unless noted otherwise. The goal of sequential disentanglement is to extract an alternative representation of $\mathbf{x}^{1:V}$ via a single static factor \mathbf{s} and multiple dynamic factors $\mathbf{d}^{1:V}$. Note that \mathbf{s} is shared across the sequence.

We can formalize the sequential disentanglement problem as a *generative task*, where every sequence $\mathbf{x}^{1:V}$ from the data space \mathcal{X} is conditioned on some $\mathbf{z}^{1:V}$ from a latent space \mathcal{Z} . We aim to maximize the probability of each sequence under the entire generative process

$$p(\mathbf{x}^{1:V}) = \int_{\mathcal{Z}} p(\mathbf{x}^{1:V} | \mathbf{z}^{1:V}) p(\mathbf{z}^{1:V}) d\mathbf{z}^{1:V}, \quad (10)$$

where $\mathbf{z}^{1:V} := (\mathbf{s}, \mathbf{d}^{1:V})$. One of the main challenges with directly maximizing Eq. (10) is that the latent space \mathcal{Z} is too large to practically integrate over. Instead, a separate distribution, denoted here as $q(\mathbf{z}^{1:V} | \mathbf{x}^{1:V})$, is used to narrow search to be only over $\mathbf{z}^{1:V}$ associated with sequences from the dataset \mathcal{D} . Importantly, the distributions $p(\mathbf{x}^{1:V} | \mathbf{z}^{1:V})$ and $q(\mathbf{z}^{1:V} | \mathbf{x}^{1:V})$ take the form of a decoder and an encoder in practice, suggesting the development of *autoencoder* sequential disentanglement models [11]. The above $p(\mathbf{x}^{1:V} | \mathbf{z}^{1:V})$ and $q(\mathbf{z}^{1:V} | \mathbf{x}^{1:V})$ are denoted by $p_{T0}(\mathbf{x}_0^\tau | \mathbf{x}_T^\tau, \mathbf{s}_0, \mathbf{d}_0^\tau)$ and $p(\mathbf{x}_t^{1:V}, \mathbf{s}_0, \mathbf{d}_0^{1:V} | \mathbf{x}_0^{1:V})$, respectively, in Eq. 1 and Eq. 2.

B.2 High-Resolution Disentangled Sequential Diffusion Autoencoder

In addition to transitioning to real-world data, our goal is to manage high-resolution data for unsupervised sequential disentanglement, for the first time. Drawing inspiration from [3], we incorporate perceptual image compression, which combines an autoencoder with a perceptual loss [58] and a patch-based adversarial objective [59, 60, 61]. Specifically, we explore two main variants of the autoencoder. The first variant applies a small Kullback–Leibler penalty to encourage the learned latent space to approximate a standard normal distribution, similar to a VAE [21, 62]. The second variant integrates a vector quantization layer [63, 64] within the decoder. Empirically, we find that the VQ-VAE-based model performs better when combined with our method. Given a pre-trained encoder \mathcal{E} and decoder \mathcal{D} , we can extract $\mathbf{x}_0^\tau = \mathcal{E}(\mathbf{x}_0^\tau)$, which represents a low-dimensional latent space where high-frequency, imperceptible details are abstracted away. Finally, \mathbf{x}_0^τ can be reconstructed from the latent \mathbf{x}_0^τ by applying the decoder $\mathbf{x}_0^\tau = \mathcal{D}(\mathbf{x}_0^\tau)$. The EDM formulation in Eq. 4 makes relatively strong assumptions about the mean and standard deviation of the training data. To meet these assumptions, we opt to normalize the training data globally rather than adjusting the value of σ_{data} , which could significantly affect other hyperparameters [65]. Therefore, we keep σ_{data} at its default value of 0.5 and ensure that the latents have a zero mean during dataset preprocessing. When generating sequence elements, we reverse this normalization before applying \mathcal{D} .

B.3 Prior Modeling

We model the prior static and dynamic distribution with $p_{T0}(\mathbf{s}_0, \mathbf{d}_0^{1:V} | \mathbf{s}_T, \mathbf{d}_T^{1:V})$. To sample static and dynamic factors, we train a separate latent DDIM model [66]. Then, we can extract the factors by sampling noise, and reversing the trained model. Specifically, we learn $p_{\Delta t}(\mathbf{z}_{t-1}^{1:V} | \mathbf{z}_t^{1:V})$ where $\mathbf{z}_0 = (\mathbf{s}_0, \mathbf{d}_0^{1:V})$ are the outputs of our sequential semantic encoder. The training is done by simply optimizing the $\mathcal{L}_{\text{latent}}$ with respect to DDIM’s output $\varepsilon_\phi(\cdot)$:

$$\mathcal{L}_{\text{latent}} = \sum_{t=1}^T \mathbb{E}_{\mathbf{z}^{1:V}, \varepsilon_t} [\|\varepsilon_\phi(\mathbf{z}_t^{1:V}, t) - \varepsilon_t\|] \quad (11)$$

where $\varepsilon_t \in \mathbb{R}^{dV+s} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, V is the sequence length, s, d are the static and dynamic factors dimensions respectively. Additionally, $\mathbf{z}_t^{1:V}$ is the noise version of \mathbf{z}_t as described in [66]. For designing the architecture of our latent model, we follow [16] and it is based on 10 MLP layers. Our network architecture and hyperparameters are provided in Tab. 8.

B.4 Reverse processes

The detailed reverse sampling algorithm is provided in Alg. 1. We follow [20] sampling techniques, however, each step in our reverse process is conditioned on the latent static and dynamic factors extracted by our sequential semantic encoder. As in [16], we observe that auto-encoding is improved significantly when using the stochastic encoding technique. Since we have a different reverse process, we provide the algorithm for stochastic encoding for our modeling in Alg. 2. Finally, when performing conditional swapping, we observe that performing stochastic encoding on the sample from which we borrow the dynamics and using it as an input to Alg. 1, improves the results empirically. That is, given two sample videos $\mathbf{x}, \hat{\mathbf{x}} \sim p_0$, to create a new sample $\tilde{\mathbf{x}}$, conditioned on the static factor of \mathbf{x} and dynamic features of $\hat{\mathbf{x}}$, we use the stochastic encoding of $\hat{\mathbf{x}}$ in Alg. 1.

C Hyper-parameters

The hyperparameters used in our autoencoder are listed in Tab. 6 and Tab. 7, detailing the configurations for each dataset: MUG, TaiChi-HD, VoxCeleb, CelebV-HQ, TIMIT, LibriSpeech, PhysioNet,

Algorithm 1 Conditioned Stochastic Sampler with $\sigma(t) = t$ and $s(t) = 1$.

```

1: procedure CONDITIONEDSTOCHASTICSAMPLER( $D_\theta$ ,  $t_i \in \{0, \dots, N\}$ ,  $\gamma_i \in \{0, \dots, N-1\}$ ,  $\mathbf{z}_0^{1:V}$ ,  $\mathbf{x}_0^{1:V}$ ,  $S_{\text{noise}}^2$ )
2:   if  $\mathbf{x}_0^{1:V} \neq \text{None}$  then
3:      $\mathbf{x}_N^{1:V} \leftarrow$  Algorithm 2 output
4:   else
5:     sample  $\mathbf{x}_N^{1:V} \sim \mathcal{N}(\mathbf{0}, t_N^2 \mathbf{I})$ 
6:     for  $i \in \{N, \dots, 1\}$  do
7:       sample  $\epsilon_i \sim \mathcal{N}(\mathbf{0}, S_{\text{noise}}^2 \mathbf{I})$ 
8:        $\hat{t}_i \leftarrow t_i + \gamma_i t_i$ 
9:        $\hat{\mathbf{x}}_i^\tau \leftarrow \mathbf{x}_i^\tau + \sqrt{\hat{t}_i^2 - t_i^2} \epsilon_i$ 
10:       $\mathbf{d}_i \leftarrow (\hat{\mathbf{x}}_i^\tau - D_\theta(\hat{\mathbf{x}}_i^\tau, \mathbf{z}_0^\tau; \hat{t}_i)) / \hat{t}_i$ 
11:       $\mathbf{x}_{i-1}^\tau \leftarrow \mathbf{x}_i^\tau + (t_{i-1} - \hat{t}_i) \mathbf{d}_i$ 
12:      if  $t_{i-1} \neq 0$  then
13:         $\mathbf{d}'_i \leftarrow (\mathbf{x}_{i-1}^\tau - D_\theta(\mathbf{x}_{i-1}^\tau, \mathbf{z}_0^\tau; t_{i-1})) / t_{i-1}$ 
14:         $\mathbf{x}_{i-1}^\tau \leftarrow \hat{\mathbf{x}}_i^\tau + (t_{i-1} - \hat{t}_i) (\frac{1}{2} \mathbf{d}_i + \frac{1}{2} \mathbf{d}'_i)$ 
15:    return  $\mathbf{x}_0$ 

```

$\triangleright \gamma_i = \begin{cases} \min(\frac{S_{\text{churn}}}{N}, \sqrt{2}-1) & \text{if } t_i \in [S_{\text{min}}, S_{\text{max}}] \\ 0 & \text{otherwise} \end{cases}$
 \triangleright Select temporarily increased noise level \hat{t}_i
 \triangleright Add new noise to move from t_i to \hat{t}_i
 \triangleright Evaluate $d\mathbf{x}/dt$ at \hat{t}_i
 \triangleright Take Euler step from t_i to t_{i-1}
 \triangleright Apply 2nd order correction

Algorithm 2 Stochastic Encoding with $\sigma(t) = t$ and $s(t) = 1$.

```

1: procedure STOCHASTICENSENCODER( $D_\theta$ ,  $t_i \in \{0, \dots, N\}$ ,  $\gamma_i \in \{0, \dots, N-1\}$ ,  $\mathbf{x}_0^{1:V}$ ,  $\mathbf{z}_0^{1:V}$ )
2:   sample  $x_0 \sim \mathcal{N}(\mathbf{0}, t_0^2 \mathbf{I})$ 
3:   for  $i \in \{0, \dots, N-1\}$  do
4:     sample  $\epsilon_i \sim \mathcal{N}(\mathbf{0}, S_{\text{noise}}^2 \mathbf{I})$ 
5:      $\hat{t}_i \leftarrow t_i + \gamma_i t_i$ 
6:      $\hat{\mathbf{x}}_i \leftarrow \mathbf{x}_i + \sqrt{\hat{t}_i^2 - t_i^2} \epsilon_i$ 
7:      $\mathbf{d}_i \leftarrow (\hat{\mathbf{x}}_i^\tau - D_\theta(\hat{\mathbf{x}}_i^\tau, \mathbf{z}_0^\tau; t_i)) / t_i$ 
8:      $\mathbf{x}_{i+1}^\tau \leftarrow \mathbf{x}_i^\tau + (t_{i+1} - t_i) \mathbf{d}_i$ 
9:     if  $t_{i+1} \neq \sigma_{\text{max}}$  then
10:       $\mathbf{d}'_i \leftarrow (\mathbf{x}_{i+1}^\tau - D_\theta(\mathbf{x}_{i+1}^\tau, \mathbf{z}_0^\tau; t_{i+1})) / t_{i+1}$ 
11:       $\mathbf{x}_{i+1}^\tau \leftarrow \hat{\mathbf{x}}_i^\tau + (t_{i+1} - \hat{t}_i) (\frac{1}{2} \mathbf{d}_i + \frac{1}{2} \mathbf{d}'_i)$ 
12:   return  $\mathbf{x}_N^{1:V}$ 

```

$\triangleright \gamma_i = \begin{cases} \min(\frac{S_{\text{churn}}}{N}, \sqrt{2}-1) & \text{if } t_i \in [S_{\text{min}}, S_{\text{max}}] \\ 0 & \text{otherwise} \end{cases}$
 \triangleright Select temporarily increased noise level \hat{t}_i
 \triangleright Add new noise to move from t_i to \hat{t}_i
 \triangleright Evaluate $d\mathbf{x}^\tau/dt$ at t_i
 \triangleright Take Euler step from t_i to t_{i+1}
 \triangleright Apply 2nd order correction

630 Air Quality and ETTh1. We provide the values of essential parameters such as sequence lengths,
631 batch sizes, learning rates, and the use of P_{mean} and P_{std} to manage noise disturbance during training.
632 In addition, the table specifies whether VQ-VAE was employed. Tab. 8 outlines the architecture of
633 our latent DDIM model, including batch size, number of epochs, MLP layers, hidden sizes, and the β
634 scheduler. These details are essential for understanding the model’s structure and its training process.
635 For the VQ-VAE model, we utilized the pre-trained model from [3] with hyperparameters $f = 8$,
636 $Z = 256$, and $d = 4$, which encodes a frame of size $3 \times 256 \times 256$ into a latent representation of
637 size $4 \times 32 \times 32$.

638 D Datasets

639 **MUG.** The MUG facial expression dataset, introduced by [49], contains image sequences from 52
640 subjects, each displaying six distinct facial expressions: anger, fear, disgust, happiness, sadness, and
641 surprise. Each video sequence in the dataset ranges from 50 to 160 frames. To create sequences of
642 length 15, as done in prior work [12], we randomly select 15 frames from the original sequences. We
643 then apply Haar Cascade face detection to crop the faces and resize them to 64×64 pixels, resulting
644 in sequences of $x \in \mathbb{R}^{15 \times 3 \times 64 \times 64}$. The final dataset comprises 3,429 samples. In the case of of the
645 zero shot experiments we resize the images to 256×256 pixels.

646 **TaiChi-HD.** The TaiChi-HD dataset, introduced by [18], contains videos of full human bodies
647 performing Tai Chi actions. We follow the original preprocessing steps from FOMM [18] and use
648 a 64×64 version of the dataset. The dataset comprises 3,081 video chunks with varying lengths,
649 ranging from 128 to 1,024 frames. We split the data into 90% for training and 10% for testing. To
650 create sequences of length 10, similar to the approach used for the MUG dataset, we randomly select

Table 6: Hyperparameters for Video datasets.

Dataset	MUG	TaiChi-HD	VoxCeleb	CelebV-HQ
P_{maen}	-1.2	-1.2	-0.4	-0.4
P_{std}	1.2	1.2	1.0	1.0
NFE	71	63	63	63
lr	1e-4	1e-4	1e-4	1e-4
bsz	8	16	16	16
#Epoch	1600	40	100	450
Dataset repeats	1	150	1	1
s dim	256	512	512	1024
d dim	64	64	12	16
hidden dim	128	1024	1024	1024
Base channels	64	64	192	192
Channel multipliers			[1, 2, 2, 2]	
Attention placement			[2]	
Encoder base ch	64	64	192	192
Encoder ch. mult.			[1, 2, 2, 2]	
Enc. attn. placement			[2]	
Input size	$3 \times 64 \times 64$	$3 \times 64 \times 64$	$3 \times 256 \times 256$	$3 \times 256 \times 256$
Seq len	15	10	10	10
Optimizer		AdamW (weight decay= $1e-5$)		
Backbone		Unet		
GPU	1 RTX 4090		3 RTX 4090	

Table 7: Hyperparameters for audio and TS.

Dataset	TIMIT	LibriSpeech	Physionet	Airq	ETTH
P_{maen}	−0.4	−0.4	−0.4	−0.4	−0.4
P_{std}	1.0	1.0	1.0	1.0	1.0
NFE	63	63	63	63	63
lr	1e−4	1e−3	5e−5	1e−4	1e−4
bsz	128	128	30	10	10
#Epoch	750	200	200	200	200
s dim	32	32	24	16	16
d dim	4	2	2	4	4
hidden dim	128	256	96	512	512
Base channels	256	64	256	256	128
Channel multipliers			[4, 4, 4, 4]		
Attention placement			None		
Encoder base ch	128	128	96	128	256
Encoder ch. mult.			[4, 4, 4, 4]		
Enc. attn. placement			None		
Input size	80	80	10	10	6
Seq len	68	68	80	672	672
Optimizer		AdamW (weight decay= 1e−5)			
Backbone		MLP			
GPU		1 RTX 4090			

10 frames from the original sequences. The resulting sequences are resized to 64×64 pixels, forming $x \in \mathbb{R}^{10 \times 3 \times 64 \times 64}$.

VoxCeleb. The VoxCeleb dataset [47] is a collection of face videos extracted from YouTube. We used the preprocessing steps from [67], where faces are extracted, and the videos are processed at 25/6 fps. The dataset comprises 22,496 videos and 153,516 video chunks. We used the verification split, which includes 1,211 speakers in the training set and 40 different speakers in the test set, resulting in 148,642 video chunks for training and 4,874 for testing. To create sequences of length 10, we

Table 8: Network architecture of our latent DDIM.

Parameter	MUG	TaiChi-HD	VoxCeleb	CelebV-HQ
Batch size	128	128	128	128
#Epoch	500	500	200	1000
MLP layers (N)			10	
MLP hidden size	1216	5008	2528	4736
β scheduler			Linear	
Learning rate			$1e-4$	
Optimizer		AdamW (weight decay= $1e-5$)		
Train Diff T			1000	
Diffusion loss		L2 loss with noise prediction ϵ		
GPU			1 RTX 4090	

randomly select 10 frames from the original sequences. The videos are processed at a resolution of 256×256 resulting in sequences represented as $x \in \mathbb{R}^{10 \times 3 \times 256 \times 256}$.

CelebV-HQ. The CelebV-HQ dataset [48] is a large-scale collection of high-quality video clips featuring faces, extracted from various online sources. The dataset consists of 35,666 video clips involving 15,653 identities, with each clip manually labeled with 83 facial attributes, including 40 appearance attributes, 35 action attributes, and 8 emotion attributes. The videos were initially processed at a resolution of 512×512 . We then used [40] to crop the facial regions, resulting in videos at a 256×256 resolution. To create sequences of length 10, we randomly selected 10 frames from the original sequences, producing sequences represented as $x \in \mathbb{R}^{10 \times 3 \times 256 \times 256}$.

TIMIT. The TIMIT dataset, introduced by [50], is a collection of read speech designed for acoustic-phonetic research and other speech-related tasks. It contains 6300 utterances, totaling approximately 5.4 hours of audio recordings, from 630 speakers (both men and women). Each speaker contributes 10 sentences, providing a diverse and comprehensive pool of speech data. To pre-process the data we use mel-spectrogram feature extraction with 8.5ms frame shift applied to the audio. Subsequently, segments of 580ms duration, equivalent to 68 frames, are sampled from the audio and treated as independent samples.

LibriSpeech. The LibriSpeech dataset [51] is a corpus of read English speech derived from audiobooks, containing 1,000 hours of speech sampled at 16 kHz. For our training, we used the `train-clean-360` subset, which consists of 363.6 hours of speech from 921 speakers. As validation and test sets, we use `dev-clean` and `test-clean`, each containing 5.4 hours of speech from 40 unique speakers, where there is no identity overlap across all subsets. For pre-processing, we extract mel-spectrogram features with an 8.5 ms frame shift applied to the audio. We then sample segments of 580 ms duration (equivalent to 68 frames) from the audio, treating them as independent samples.

PhysioNet. The PhysioNet ICU dataset [54] consists of medical time series data collected from 12,000 adult patients admitted to the Intensive Care Unit (ICU). This dataset includes time-dependent measurements such as physiological signals, laboratory results, and relevant patient demographics like age and reasons for ICU admission. Additionally, labels indicating in-hospital mortality events are included. Our preprocessing procedures follow the guidelines provided in [52].

Air Quality. The UCI Beijing Multi-site Air Quality dataset [55] comprises hourly records of air pollution levels, collected over a four-year period from March 1, 2013, to February 28, 2017, across 12 monitoring sites in Beijing. Meteorological data from nearby weather stations of the China Meteorological Administration is also included. Our approach to data preprocessing, as described in [52], involves segmenting the data based on different monitoring locations and months of the year.

ETTh1. The ETTh1 dataset is a subset of the Electricity Transformer Temperature (ETT) dataset, containing hourly data over a two-year period from two counties in China. The dataset is focused on Long Sequence time series Forecasting (LSTF) of transformer oil temperatures. Each data point

consists of the target value (oil temperature) and six power load features. The dataset is divided into training, validation, and test sets, with a 12/4/4-month split.

E Metrics

Average Keypoint Distance (AKD). To evaluate whether the motion in the reconstructed video is preserved, we utilize pre-trained third-party keypoint detectors on the TaiChi-HD, VoxCeleb, CelebV-HQ, and MUG datasets. For the VoxCeleb, CelebV-HQ and MUG datasets, we employ the facial landmark detector from [68], whereas for the TaiChi-HD dataset, we use the human-pose estimator from [69]. Keypoints are computed independently for each frame. AKD is calculated by averaging the L_1 distance between the detected keypoints in the ground truth and the generated video. The TaiChi-HD and MUG datasets are evaluated at a resolution of 64×64 pixels, and the VoxCeleb and CelebV-HQ datasets at 256×256 pixels. If the model output is at a lower resolution, it is interpolated to 256×256 pixels for evaluation.

Average Euclidean Distance (AED). To assess whether the identity in the reconstructed video is preserved, we use the Average Euclidean Distance (AED) metric. AED is calculated by measuring the Euclidean distance between the feature representations of the ground truth and the generated video frames. We selected the feature embedding following the example set in [18]. For the VoxCeleb, CelebV-HQ, and MUG datasets, we use a VGG-FACE for facial identification using the framework of [70], whereas for TaiChi-HD, we use a network trained for person re-identification [71]. TaiChi-HD and MUG are evaluated at a resolution of 64×64 pixels, and VoxCeleb and CelebV-HQ at 256×256 pixels.

To ensure fairness when measuring AED and AKD, we created a predefined dataset of example pairs, ensuring that all models are evaluated on the exact same set of pairs. This is important because when measuring quantitative metrics, the results may vary depending on the dynamics swapped between two subjects, as e.g., the key points in AKD in the original video are influenced by the identity of the person. To address this issue, we establish a fixed set of pairs for a consistent comparison across all methods.

Accuracy (Acc). As in [14], we used this metric for the MUG dataset to evaluate a model’s ability to preserve fixed features while generating others. For example, dynamic features are frozen while static features are sampled. Accuracy is computed using a pre-trained classifier, referred to as the “judge”, which is trained on the same training set as the model and tested on the same test set. For the MUG dataset, the classifier checks that the facial expression remains unchanged during the sampling of static features.

Inception Score (IS). The Inception Score is a metric used to evaluate the performance of the model generation. First, we apply the judge, to all generated videos $x_0^{1:V}$, obtaining the conditional predicted label distribution $p(y|x_0^{1:V})$. Next, we compute $p(y)$, the marginal predicted label distribution, and calculate the KL-divergence $\text{KL}[p(y|x_0^{1:V}) \| p(y)]$. Finally, the Inception Score is computed as $\text{IS} = \exp(\mathbb{E}_x \text{KL}[p(y|x_0^{1:V}) \| p(y)])$. We use this metric evaluate our results on MUG dataset.

Inter-Entropy ($H(y|x_0^{1:V})$). This metric reflects the confidence of the judge in its label predictions, with lower inter-entropy indicating higher confidence. It is calculated by passing k generated sequences $\{x_0^{1:V}\}^{1:k}$ into the judge and computing the average entropy of the predicted label distributions: $\frac{1}{k} \sum_{i=1}^k H(p(y|\{x_0^{1:V}\}^i))$. We use this metric evaluate our results on MUG dataset.

Intra-Entropy ($H(y)$). This metric measures the diversity of the generated sequences, where a higher intra-entropy score indicates greater diversity. It is computed by sampling from the learned prior distribution $p(y)$ and then applying the judge to the predicted labels y . We use this metric to evaluate our results on the MUG dataset.

EER. Equal Error Rate (EER) metric is widely employed in speaker verification tasks. The EER represents the point at which the false positive rate equals the false negative rate, offering a balanced measure of performance in speaker recognition. This metric, commonly applied to the TIMIT dataset, provides a robust evaluation of the model’s ability to disentangle features relevant to speaker identity.

743 **AUPRC.** The Area Under the Precision-Recall Curve (AUPRC) is a metric that evaluates the
744 balance between precision and recall by measuring the area beneath their curve. A higher AUPRC
745 reflects superior model performance, with values nearing 1 being optimal, indicating both high
746 precision and recall.

747 **AUROC.** The Area Under the Receiver Operating Characteristic Curve (AUROC) measures the
748 trade-off between true positive rate (TPR) and false positive rate (FPR), quantifying the area under
749 the curve of these rates. A higher AUROC signifies better performance, with values close to 1 being
750 desirable, representing a model that distinguishes well between positive and negative classes.

751 **MAE.** Mean Absolute Error (MAE) calculates the average magnitude of errors between predicted
752 and observed values, offering a simple and intuitive measure of model accuracy. As it computes the
753 average absolute difference between predicted and actual values, MAE is resistant to outliers and
754 provides a clear indication of the model’s prediction precision.

755 **DNSMOS.** Deep Noise Suppression Mean Opinion Score (DNSMOS [72]) is a neural network-
756 based metric introduced to estimate the perceptual quality of speech processed by noise suppression
757 algorithms. Trained to predict human Mean Opinion Scores (MOS), DNSMOS provides a no-
758 reference quality assessment that correlates strongly with subjective human judgments. It evaluates
759 both the speech quality and the effectiveness of noise reduction, offering a comprehensive measure of
760 audio clarity and intelligibility. This metric is especially useful in evaluating real-world performance
761 of speech enhancement systems without the need for costly and time-consuming human listening
762 tests.




Figure 5: Rows A and B are two inputs from the test set. Row C shows a dynamic swap example, using the static of A and dynamics of B. In row D we extract the same person from A, but with the dynamics as labeled in B. Finally, in row E, we extract the same person from A with the dynamics that are predicted by the classifier.



Figure 6: Rows A and B are two inputs from the test set. Row C shows a dynamic swap example, using the static of A and dynamics of B. In row D we extract the same person from A, but with the dynamics as labeled in B. Finally, in row E, we extract the same person from A with the dynamics that are predicted by the classifier.

Table 9: Judge benchmark disentanglement metrics on MUG.

Method	Acc \uparrow	IS \uparrow	MUG		Reconstruction (MSE) \downarrow
			$H(y x)\downarrow$	$H(y)\uparrow$	
MoCoGAN	63.12%	4.332	0.183	1.721	—
DSVAE	54.29%	3.608	0.374	1.657	—
R-WAE	71.25%	5.149	0.131	1.771	—
S3VAE	70.51%	5.136	0.135	1.760	—
SKD	77.45%	5.569	0.052	1.769	—
C-DSVAE	81.16%	5.341	0.092	1.775	—
SPYL	85.71%	5.548	0.066	1.779	1.311e−3
DBSE	86.90%	5.598	0.041	1.782	1.286e−3
Ours	81.15%	5.382	0.090	1.773	2.669e−7

F MUG and Judge Metric Analysis

While our results show significant improvement over previous methods on VoxCeleb [47], CelebV-HQ [48], and TaiChi-HD [18], both in terms of disentanglement and reconstruction, our performance on MUG [49] is only on par with the state-of-the-art methods. Since MUG is a labeled dataset, the traditional evaluation task involves the unconditional generation of static factors while freezing the dynamics, resulting in altering the appearance of the person. The generated samples are then evaluated using an off-the-shelf judge model (See App. E), which is a neural network trained to classify both static and dynamic factors. If the disentanglement method disentangles these factors effectively, we expect the judge to correctly identify the dynamics while outputting different predictions for the static features, since the latter were randomly sampled and should differ from the original static factor.

Surprised by our results on MUG, we investigated the failure cases to understand the limitations of our model. In particular, we examined scenarios where we freeze the dynamics and swap the static features between two samples, and then we generate the corresponding output. In Fig. 5, we show an example where the static features of the second row are swapped with those of the first row, and the resulting generation is displayed in the third row. We observe that while the dynamics from the second row are well-preserved, the generated person retains the identity of the first row. However, the classifier incorrectly predicts the dynamics for the sequence. To further investigate this, we extracted a ground-truth example of the person from the first row in the dataset expressing the expected emotion and the predicted one. In the last two rows of Fig. 5, we show the same person with predicted dynamics (fourth row) and the same person with the dynamics that the classifier predicted (fifth row). We provide another example of the same phenomenon in Fig. 6.

We observe that while the judge predicts the wrong label for our generated samples in rows C, the facial expressions of the people there align better with the actual dynamics in rows B. This suggests that the classifier is biased towards the identity when predicting dynamics, potentially forming a discrete latent space where generalization to nearby related expressions is not possible. Importantly, the judge attains $> 99\%$ accuracy on the test set. We conclude that utilizing a judge can be problematic for measuring new and unseen variations in the data. This analysis motivates us to present the AKD and AED, as detailed above in App. E.

G Additional Experiments

G.1 Dependent vs. Independent Prior Modeling

In Sec. 3, we describe our approach to prior modeling, highlighting our decision to generate latent factors dependently rather than independently, as done in previous state-of-the-art methods. Beyond being a parameter- and time-efficient choice, we empirically validate the advantages of our approach in the following experiment.

In this experiment, we compare two setups: (1) dependent generation of static and dynamic latent vectors, and (2) independent generation of these latent vectors using two latent DDIM models: one for the static vector and another for the dynamic vectors. To quantitatively assess the effectiveness of both approaches, we measure the Fréchet Video Distance (FVD) [73], a metric derived from the

well-established FID score for videos. This metric evaluates how well a generative model captures the observed data distribution, where lower scores indicate better performance.

We conduct our evaluation on the VoxCeleb dataset, training two latent models. The independent model achieves an FVD score of 75.03, whereas our dependent approach achieves a significantly lower score of **65.23**, representing a $\approx 13\%$ improvement. This result underscores the expressive advantage of modeling latent factors dependently.

G.2 Additional Analysis of DiffSDA Disentanglement Components

This section explores the impact of two key components of our method on disentanglement quality: i) the static latent factor s_0 shared across all time steps τ , and ii) the dimensionality of the dynamic latent factor d_0^τ .

To analyze these effects, we trained four models on the VoxCeleb dataset for 100 epochs, maintaining a static latent dimension of 128 while varying the size of the dynamic latent factor and whether the static latent factor was shared or not. The models were evaluated using our conditional swapping protocol and a verification metric based on the VGG-FACE framework proposed in [70]. Specifically, we assessed identity consistency by freezing the static factor and swapping the dynamic factor, with the verification score representing the percentage of cases where identity was correctly preserved across frames.

As shown in Tab.10, our results indicate that the optimal performance (first row of the table) is achieved when d_0^τ has a smaller dimensionality, and the static factor is shared. Other configurations reveal significant trade-offs: increasing d_0^τ dimensionality results in higher AED scores but reduced verification accuracy, indicating weaker disentanglement of the static factor. Similarly, when s_0 is not shared, the AKD score degrades significantly, suggesting ineffective disentanglement of the dynamic factor. These findings underscore the importance of both (i) and (ii) in achieving robust sequential disentanglement.

Table 10: Disentanglement effect of VoxCeleb dataset

d_0^τ size	s shared?	Verification ACC \uparrow (Static Frozen)	AED \downarrow (static frozen)	AKD \downarrow (dynamics frozen)
16	✓	64.36%	0.925	2.882
128	✓	18.03%	1.054	2.077
16	✗	56.75%	0.898	12.64
128	✗	48.41%	0.980	12.28

G.3 Speech Quality and Reconstruction Comparison

This section discusses the results of speech reconstruction and quality evaluation presented in table 11 on the LibriSpeech dataset. We compare the reconstruction performance using the Mean Squared Error (MSE) on the spectrograms and assess speech quality using the Deep Noise Suppression Mean Opinion Score (DNSMOS) [72]. The DNSMOS metric has a maximum score of 5, but the original (reference) dataset achieves a score of 3.9, as shown in the REF row of the table. As can be seen in the table, our model outperforms all comparable methods, achieving the lowest MSE and the highest DNSMOS among the evaluated approaches.

Table 11: Disentanglement and generation quality metrics on Libri Speech. For generation quality, we report MSE on the spectrogram and Deep Noise Suppression Mean Opinion Score (DNSMOS).

	Method	MSE \downarrow	DNSMOS \uparrow
Libri Speech	REF	--	3.9
	DSVAE	5.53e-2	3.13
	SPYL	4.40e-1	2.21
	DBSE	6.72e-3	2.88
	Ours	1.83e-4	3.41

G.4 Generative Quality Compression

This section discusses the generative quality results shown in Table 12, evaluated using the Fréchet Video Distance (FVD) on the VoxCeleb dataset. We generated the same number of samples as in the test set and computed the FVD score against the test set. This process was repeated five times for each model using different five different seeds to obtain a robust estimate. We report the mean FVD along with the standard deviation. The results demonstrate that our model outperforms existing state-of-the-art sequential disentanglement models in the video generation task.

Table 12: Fréchet Video Distance (FVD) results on VoxCeleb dataset to assess video generation quality. All experiments were conducted across five different random seeds to ensure robustness and account for variability in generation.

Model	FVD↓
SPYL	582.28 ± 1.15
DBSE	1076.44 ± 2.22
Ours	65.23 ± 0.81

H Additional Results

H.1 Reconstruction Results

In Figs. 7 to 10, we present several qualitative reconstruction examples across all datasets.

H.2 Additional results: conditional swap

In what follows, we present more results for the conditional swapping experiment from the main text (Sec. 4.1). In each figure, the first two rows show the original sequences (real videos). The third and fourth rows are the results of the conditional swap where we change the dynamic and static factors, respectively. We show our results for all datasets in Figs. 11 to 14.

H.3 Additional results: unconditional swap

In addition to the conditional and zero-shot tasks considered above, we can also perform such tasks in an unconditional manner. Specifically, given a real sequence $\mathbf{x}^{1:V}$ with its factors $(\mathbf{s}, \mathbf{d}^{1:V})$, we can unconditionally sample new $(\hat{\mathbf{s}}, \hat{\mathbf{d}}^{1:V})$ using our separate DDIM model (see Sec. 3). We then reconstruct the static swap $(\hat{\mathbf{s}}, \mathbf{d}^{1:V})$ and the dynamic swap $(\mathbf{s}, \hat{\mathbf{d}}^{1:V})$ similarly as described above. In Fig. 15, we present unconditional swap results on CelebV-HQ (left), VoxCeleb (middle), and TaiChi-HD (right). The middle rows represent the original sequences, whereas the top and bottom rows demonstrate dynamic and static swaps, respectively. Across all datasets and swap settings, our approach succeeds in modifying the swapped features while preserving the frozen factors, either in the static or in the dynamic examples. In addition, we also present more results where each figure is composed of separate panels. In each panel, the middle row represents the original sequence. In the top row, we sample new dynamic factors and freeze the static factor. In the bottom row below, we sample a new static factor and freeze the dynamics. We show our results on all datasets in Figs. 16 to 19.

H.4 Additional results: zero-shot disentanglement

Here we extend the results from Sec. 4.2. We provide additional examples of conditional swapping when the model is trained on one dataset and evaluated on another dataset, unseen during training. Specifically, in Fig. 20, we show examples where the model is trained on VoxCeleb and tested on MUG. Additionally, in Fig. 21, the model is trained on VoxCeleb and tested on CelebV-HQ. Finally, in Fig. 22, the model is trained on CelebV-HQ and tested on VoxCeleb.




Figure 7: Reconstruction results of CelebV-HQ (256×256). The first row for each pair is the original video and the second row is its reconstruction.

868 H.5 Additional results: multifactor disentanglement

869 In this section, we present more examples for traversing the latent space, separately for the static and
 870 dynamic factors. For static factors, we show in Figs. 24 to 35. There, we find different factors of
 871 variation such as Male to Female, younger to older, brighter and darker hair color, and more. Each
 872 row in the figure is a video, and the different columns represent the traversal in α values (see Eq. 6).
 873 In addition, we present full examples of dynamic factor traversal in Figs. 36 to 47, demonstrating
 874 various factors of variation. Among the factors are facial expressions, camera angles, head rotations,
 875 eyes and mouth control, etc.




Figure 8: Reconstruction results of VoxCeleb (256×256). The first row for each pair is the original video and the second row is its reconstruction.




Figure 9: Reconstruction results of TaiChi-HD. The first row for each pair is the original video and the second row is its reconstruction.




Figure 10: Reconstruction results of MUG. The first row for each pair is the original video and the second row is its reconstruction.




Figure 11: Each panel contains a pair of original videos from CelebV-HQ (Real videos), and a pair of conditional swapping of the dynamic and static factors (Swapped videos).

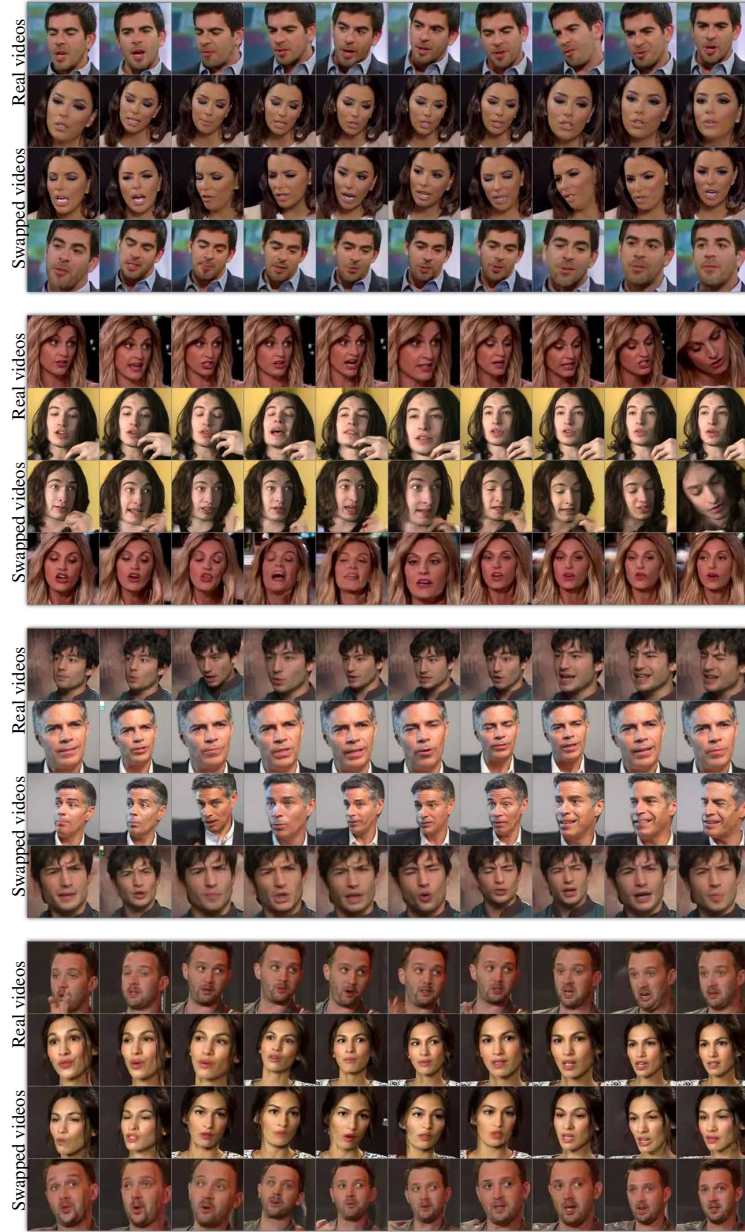


Figure 12: Each panel contains a pair of original videos from VoxCeleb (Real videos), and a pair of conditional swapping of the dynamic and static factors (Swapped videos).




Figure 13: Each panel contains a pair of original videos from TaiChi-HD (Real videos), and a pair of conditional swapping of the dynamic and static factors (Swapped videos).




Figure 14: Each panel contains a pair of original videos from MUG (Real videos), and a pair of conditional swapping of the dynamic and static factors (Swapped videos).



Figure 15: Unconditional dynamic (top) and static (bottom) swap results on CelebV-HQ (left), VoxCeleb (middle), and TaiChi-HD (right).




Figure 16: CelebV-HQ unconditional swapping. The middle row represents the original video (real), the row above shows a dynamic swap (dynamics), and the row below shows a static swap (static).



Figure 17: VoxCeleb unconditional swapping. The middle row represents the original video (real), the row above shows a dynamic swap (dynamics), and the row below shows a static swap (static).




Figure 18: TaiChi-HD unconditional swapping. The middle row represents the original video (real), the row above shows a dynamic swap (dynamics), and the row below shows a static swap (static).




Figure 19: MUG unconditional swapping. The middle row represents the original video (real), the row above shows a dynamic swap (dynamics), and the row below shows a static swap (static).




Figure 20: Each panel contains in its first and second rows a pair of real videos from VoxCeleb and MUG, respectively. We perform conditional swapping using a model that was trained on VoxCeleb, but we zero-shot swap the dynamic and static factors of a MUG example (Swapped videos).




Figure 21: Each panel contains in its first and second rows a pair of real videos from VoxCeleb and CelebV-HQ. We perform conditional swapping using a model that was trained on VoxCeleb, but we zero-shot swap the dynamic and static factors of a CelebV-HQ example (Swapped videos).




Figure 22: Each panel contains in its first and second rows a pair of real videos from CelebV-HQ and VoxCeleb. We perform conditional swapping using a model that was trained on CelebV-HQ, but we zero-shot swap the dynamic and static factors of a VoxCeleb example (Swapped videos).




Figure 23: Traversing the latent space of DiffSDA via PCA reveals multiple dynamic variations on CelebV-HQ, including surprised and serious expressions, and different head orientations.



Figure 24: Traversing between Male appearances and Female appearances.



Figure 25: Traversing over a darker hair factor.



Figure 26: Traversing between sharper and blurry videos.



Figure 27: Traversing over a brighter hair factor.



Figure 28: Traversing between younger and older appearances.



Figure 29: Traversing over skin color variations.



Figure 30: Traversing between Male appearances and Female appearances.




Figure 31: Traversing over a darker hair factor.




Figure 32: Traversing between sharper and blurry videos.




Figure 33: Traversing over a brighter hair factor.




Figure 34: Traversing between younger and older appearances.



Figure 35: Traversing over skin color variations.



Figure 36: Traversing a head rotation factor.



Figure 37: Traversing over head angles.



Figure 38: Traversing over up and down rotations.



Figure 39: Traversing over facial expressions.



Figure 40: Traversing over mouth openness factor.



Figure 41: Traversing over eyes openness factor.



Figure 42: Traversing over a head rotation factor.



Figure 43: Traversing over various head angles.



Figure 44: Traversing over up and down head rotations.



Figure 45: Traversing over facial expressions.



Figure 46: Traversing over mouth openness factor.



Figure 47: Traversing over eyes openness factor.

NeurIPS Paper Checklist

The checklist is designed to encourage best practices for responsible machine learning research, addressing issues of reproducibility, transparency, research ethics, and societal impact. Do not remove the checklist: **The papers not including the checklist will be desk rejected.** The checklist should follow the references and follow the (optional) supplemental material. The checklist does NOT count towards the page limit.

Please read the checklist guidelines carefully for information on how to answer these questions. For each question in the checklist:

- You should answer [Yes], [No], or [NA].
- [NA] means either that the question is Not Applicable for that particular paper or the relevant information is Not Available.
- Please provide a short (1–2 sentence) justification right after your answer (even for NA).

The checklist answers are an integral part of your paper submission. They are visible to the reviewers, area chairs, senior area chairs, and ethics reviewers. You will be asked to also include it (after eventual revisions) with the final version of your paper, and its final version will be published with the paper.

The reviewers of your paper will be asked to use the checklist as one of the factors in their evaluation. While "[Yes]" is generally preferable to "[No]", it is perfectly acceptable to answer "[No]" provided a proper justification is given (e.g., "error bars are not reported because it would be too computationally expensive" or "we were unable to find the license for the dataset we used"). In general, answering "[No]" or "[NA]" is not grounds for rejection. While the questions are phrased in a binary way, we acknowledge that the true answer is often more nuanced, so please just use your best judgment and write a justification to elaborate. All supporting evidence can appear either in the main paper or the supplemental material, provided in appendix. If you answer [Yes] to a question, in the justification please point to the section(s) where related material for the question can be found.

IMPORTANT, please:

- Delete this instruction block, but keep the section heading “NeurIPS Paper Checklist”.
- Keep the checklist subsection headings, questions/answers and guidelines below.
- Do not modify the questions and only use the provided macros for your answers.

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope?

Answer: [Yes]

Justification: We believe that the abstract claims are reflected in the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: See the conclusions section.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: App A and B

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: See App C for hyperparameter App D for dataset information, and section 3 for the modeling.

Guidelines:

- The answer NA means that the paper does not include experiments.

- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: Code will be provided upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).

- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: See App C and D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: See table 5.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [\[Yes\]](#)

Justification: See Tables 6, 7 and 8.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: No societal impact of the work performed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [No]

Justification: The paper poses no such risks (NA).

Guidelines:

- The answer NA means that the paper poses no such risks.

- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [\[Yes\]](#)

Justification: Citation are present for relevant models and datasets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [\[Yes\]](#)

Justification: See App C and D.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [\[NA\]](#)

Justification: Crowd-sourcing experiments and research with human subjects irrelevant for this paper.

1184 Guidelines:

1185 • The answer NA means that the paper does not involve crowdsourcing nor research with

1186 human subjects.

1187 • Including this information in the supplemental material is fine, but if the main contribu-

1188 tion of the paper involves human subjects, then as much detail as possible should be

1189 included in the main paper.

1190 • According to the NeurIPS Code of Ethics, workers involved in data collection, curation,

1191 or other labor should be paid at least the minimum wage in the country of the data

1192 collector.

1193 **15. Institutional review board (IRB) approvals or equivalent for research with human**

1194 **subjects**

1195 Question: Does the paper describe potential risks incurred by study participants, whether

1196 such risks were disclosed to the subjects, and whether Institutional Review Board (IRB)

1197 approvals (or an equivalent approval/review based on the requirements of your country or

1198 institution) were obtained?

1199 Answer: [NA]

1200 Justification: No relevant risks identified.

1201 Guidelines:

1202 • The answer NA means that the paper does not involve crowdsourcing nor research with

1203 human subjects.

1204 • Depending on the country in which research is conducted, IRB approval (or equivalent)

1205 may be required for any human subjects research. If you obtained IRB approval, you

1206 should clearly state this in the paper.

1207 • We recognize that the procedures for this may vary significantly between institutions

1208 and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the

1209 guidelines for their institution.

1210 • For initial submissions, do not include any information that would break anonymity (if

1211 applicable), such as the institution conducting the review.

1212 **16. Declaration of LLM usage**

1213 Question: Does the paper describe the usage of LLMs if it is an important, original, or

1214 non-standard component of the core methods in this research? Note that if the LLM is used

1215 only for writing, editing, or formatting purposes and does not impact the core methodology,

1216 scientific rigorousness, or originality of the research, declaration is not required.

1217 Answer: [NA]

1218 Justification: LLM used only for writing, editing and formatting purposes.

1219 Guidelines:

1220 • The answer NA means that the core method development in this research does not

1221 involve LLMs as any important, original, or non-standard components.

1222 • Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>)

1223 for what should or should not be described.