

# The Universal Normal Embedding

Chen Tasker\* Roy Betser\* Eyal Gofer\* Meir Yossef Levi Guy Gilboa  
 Viterbi Faculty of Electrical and Computer Engineering  
 Technion - Israel Institute of Technology

## Abstract

Generative models and vision encoders have largely advanced on separate tracks, optimized for different goals and grounded in different mathematical principles. Yet, they share a fundamental property: latent space Gaussianity. Generative models map Gaussian noise to images, while encoders map images to semantic embeddings whose coordinates empirically behave as Gaussian. We hypothesize that both are views of a shared latent source, the Universal Normal Embedding (UNE): an approximately Gaussian latent space from which encoder embeddings and DDIM-inverted noise arise as noisy linear projections. To test our hypothesis, we introduce NoiseZoo, a dataset of per-image latents comprising DDIM-inverted diffusion noise and matching encoder representations (CLIP, DINO). On CelebA, linear probes in both spaces yield strong, aligned attribute predictions, indicating that generative noise encodes meaningful semantics along linear directions. These directions further enable faithful, controllable edits (e.g., smile, gender, age) without architectural changes, where simple orthogonalization mitigates spurious entanglements. Taken together, our results provide empirical support for the UNE hypothesis and reveal a shared Gaussian-like latent geometry that concretely links encoding and generation. Code and data are available [here](#).

## 1. Introduction

Generative modeling has reshaped visual computing, enabling high-fidelity synthesis, reconstruction, and editing [20, 25, 29, 49]. In parallel, foundation models have learned highly semantic representations through self-supervision, where simple linear heads achieve strong classification, retrieval, and zero-shot recognition [13, 15, 45]. Together, these advances shifted vision from passive recognition to general-purpose creation and understanding, now spanning diverse visual domains [3, 44].

\*These authors contributed equally to this work.

Corresponding author: roybe@campus.technion.ac.il




Figure 1. **UNE conceptual illustration.** Different encoders (e.g., CLIP, DINO) and generative models (e.g., SD, LCM) provide different views of the same underlying Gaussian latent structure. Although trained for different objectives, their latents can be interpreted as noisy linear projections of a shared ideal Gaussian space.

Prior work reveals surprising *linearity* and even shared geometry across deep latent spaces [4, 63]. First, within *generative* families, independently trained VAEs, GANs, flows, and diffusion models can be “stitched” (i.e., their latent spaces can be linearly aligned so that codes from one model can be decoded by another) via simple linear maps between their latents [2, 4, 32, 39, 63]. Similarly, within *representation* families, vision encoders likewise “stitch” across architectures and modalities. Single-projection text-image alignment and shallow model-stitching show that independently trained encoders can operate in a shared latent space [7, 30, 36, 40, 59].

Motivated by the Platonic Representation Hypothesis and embedding-translation results [26, 27], and by identifiability showing that contrastive encoders invert the data-generating process [64], we unify the encoder and generator worlds by directly linking generative noise to encoder representations. We posit a shared, approximately Gaussian latent space, the *Universal Normal Embedding (UNE)*, from which both families arise as *noisy linear projections*.




Figure 2. **Universal Normal Embedding (UNE)**. A multivariate standard Gaussian latent space representing the encoded data distribution, in which linear directions align with semantics: classes are separable by hyperplanes, and continuous attributes (e.g., “smile”) can be edited by perturbing along a single latent direction.

UNE refers to an ideal Gaussian latent space whose linear projections approximate the latent spaces of both generative models and vision encoders (see illustration in Figure 1). In this geometry, semantic variation aligns with linear directions [12], making UNE *actionable* for linear probes and controllable edits (illustrated in Figure 2).

Evidence motivating UNE comes from both sides. Generative models sample from Gaussian priors, while encoder representations (e.g., CLIP [45], DINO [13]) empirically behave as approximately Gaussian [11, 23]. Contrastive-learning theory shows that encoders can recover the latent generative factors [64], and follow-up work establishes identifiability of encoder representations up to linear transformations [18, 47]. In parallel, large models converge toward shared latent geometry across architectures and modalities [26, 27, 34, 57]. Recent theoretical work further formalizes different regimes in which representations exhibit Gaussian behavior [5, 10]. These results suggest that encoder latents and generative noise reflect the same underlying factors. We show that these factors admit an approximately Gaussian shared latent space in practice, with encoders and generators aligning as noisy linear projections of that space.

Having established the motivation and formulation of UNE, we investigate it empirically by analyzing latent representations from multiple diffusion models and vision encoders using a unified per-image dataset. We evaluate observable consequences predicted by the hypothesis: Gaussianity of coordinates, linear separability of semantic attributes, cross-model latent alignment, and linear controllability of semantic directions. We further examine multi-view intersections of these latent spaces to study whether they preserve a consistent shared structure. Together, these evaluations suggest that encoder and generative latents behave as noisy linear views of a common, approximately

Gaussian latent source.

Our main contributions are:

- **Universal Normal Embedding (UNE)**. We formalize the UNE hypothesis of a shared, approximately Gaussian latent space linking encoders and generators, and relate it to real latents; as a proof of concept, we also explore a multi-view estimator that recovers a shared  $k$ -dimensional intersection subspace across models.
- **Semantic structure in generative noise**. We show that DDIM-inverted noise encodes rich semantics: linear probes on noise alone achieve strong attribute prediction across multiple diffusion models, closely matching foundation encoders.
- **Controllable editing via linear directions**. We enable faithful, interpretable edits by shifting along probe-derived directions in noise space, and show that a simple orthogonalization mitigates spurious entanglements, without architectural changes or fine-tuning.
- **NoiseZoo dataset**. We release *NoiseZoo*: per-image DDIM-inverted noise paired with matched encoder embeddings for real images, enabling studies of generative-semantic correspondence.

## 2. Related Work

**Latent alignment and shared geometry.** Despite architectural and objective differences, the latent spaces of VAEs [29], GANs [20], normalizing flows [48], and diffusion models [25, 56] often exhibit surprising alignment. Empirically, several works show that simple linear mappings can translate between latent spaces [2, 4, 32, 39, 63], even across models trained independently or with different dimensionalities. Other studies observe that cross-modal or cross-architecture representations remain compatible under shallow linear transforms [7, 30, 36, 40, 59]. A complementary direction seeks theoretical explanations for such

alignment. Conceptual frameworks like the Platonic Representation Hypothesis [26] and embedding translation [27] argue that diverse models converge toward a shared latent description of the scene. On the identifiability side, it was shown that InfoNCE can recover latent generative factors up to component-wise invertible transforms [64], with follow-up work tightening this to linear identifiability and cross-encoder alignment [18, 47].

However, these theoretical accounts assume a shared space without specifying its *geometry*, while empirical alignment works reveal compatibility but offer no operational mechanism for *using* the shared latent. We instead propose that this shared space is not only present but approximately *Gaussian*, making simple linear classification, semantic manipulation, and shared-space constructions natural operations that explicitly exploit its geometry.

**Gaussianity of representation spaces.** Self-supervised learning implicitly encourages isotropy: contrastive learning spreads features uniformly on the hypersphere [60], while redundancy-reduction methods decorrelate features [8, 62]. Whitening-based methods further produce Gaussianized embeddings [19], and foundation model representations exhibit approximately Gaussian statistics [9, 33]. Theory helps explain this trend: both contrastive and supervised training can recover latent factors up to linear transforms [18, 43, 47]. Additional work characterizes when representations exhibit Gaussian behavior [5, 6, 10]. Prior work has shown that multi-modal representations exhibit a modality gap and often lie in lower-dimensional, anisotropic subspaces rather than being uniformly distributed [35, 50, 53, 61]. In this work, we focus on a single modality, namely the image modality. These works, however, focus on encoder geometry only, whereas we place both encoders *and* generative models under the same approximately Gaussian latent space.

**Semantic editing in generative latents.** GANs enable editing along latent directions [22, 52], but diffusion models lack a persistent latent code. Recent approaches introduce editable subspaces [31, 58], or find directions via PCA, Jacobians or contrastive objectives [14, 17, 21]. Null-text inversion [42] and prompt-based manipulation [24] improve controllability but do not expose explicit latent semantics. Recent work exploits approximate linearity of diffusion outputs for controllable sampling [55]. Unlike these methods, we operate directly in the *noise space*, showing that it encodes semantic structure comparable to representation embeddings and enabling simple linear edits in noise space without prompt engineering or model fine-tuning.

### 3. Universal Normal Embedding (UNE)

Generative models and vision encoders share a key property: their latents exhibit approximately Gaussian structure. Yet their capabilities differ, with encoders excelling at high-

level semantic representations that support linear recognition and retrieval; in contrast, generative models carry precise pixel-level information and can synthesize or reconstruct images. For example, DDIM inversion can recover image-specific noise codes for a given diffusion model, but semantic editing in these models typically relies on external guidance (e.g., text prompts, architectural changes, or extra training) and remains limited without it. Despite these differences in objective and usage, both families access the same data distribution (e.g., natural images) and, empirically, produce Gaussianized latent variables. This complementarity motivates our central view: *encoding and generation are two related directions over a shared latent Gaussian geometry*, which we formalize as the *Universal Normal Embedding (UNE)* hypothesis.

#### Hypothesis 1: Universal Normal Embedding (UNE)

Let  $S$  be the data domain (e.g., natural images) with data sampled from a distribution  $p$ . We posit that there exists a Gaussian latent space, the *Universal Normal Embedding (UNE)*, which is  $\mathcal{N}(0, I_D)$  of unknown dimension  $D \in \mathbb{N}$  such that:

1. There exists an invertible, information-preserving mapping between  $S$  and the UNE.
2. The UNE representation is *simple* in the sense that there exists a set of semantic properties of the data, where each property is linearly separable.

### 3.1. Induced Normal Embeddings

In practice, models do not recover the full UNE for several reasons. First, their latent dimensionalities differ, often chosen heuristically to balance performance and computational cost. Second, variations in training objectives and architectures lead models to encode different aspects of the underlying information. Third, the data modalities vary: for instance, CLIP is trained on paired image-text data, whereas DINO and most generative models are not. Accordingly, both encoders and generative models realize an *Induced Normal Embedding*: a model-specific latent space that is well-approximated by a noisy linear projection of the ideal UNE. Some of the true latent structure is preserved, some dimensions may be discarded, and additional model-specific noise or redundancy may be injected.

Hence, all models are exposed to different parts of the “true” representation, varying due to different transforms and model-specific noise. An immediate consequence of our Hypotheses is that in the noiseless case ( $\epsilon_i = 0$  in Equation (1)), if  $C_i$  is invertible, any semantic property which is linearly separable in the UNE is also linearly separable in the INE. Moreover, linear separability across multiple INEs suggests a shared low-dimensional space given by their intersection, preserving separability under linear projections.

Table 1. Gaussianity measured via random 1D projections. For each model, we evaluate 5,000 projections of 250-sample subsets using Anderson-Darling (AD), D’Agostino-Pearson (DP), and Shapiro-Wilk (SW) tests (AD: lower is better ↓; DP, SW: higher is better ↑). AD%, DP% and SW% denote the fraction of projections classified as Gaussian (AD < 0.752; DP and SW  $p$ -value > 0.05). Generative models approach the theoretical 95% acceptance rate of Gaussian samples, encoders remain high, and non-Gaussian references perform substantially worse.

Model	Anderson–Darling (AD)		D’Agostino–Pearson (DP)		Shapiro–Wilk (SW)	
	Avg. AD ↓	AD % ↑	Avg. DP ↑	DP % ↑	Avg. SW ↑	SW % ↑
SD 1.5	0.387	96.00	0.504	95.16	0.491	95.02
SD 2.1	0.381	95.80	0.496	94.76	0.497	94.96
LCMv7	0.382	95.58	0.504	94.56	0.500	94.44
CLIP B16	0.451	89.50	0.430	88.88	0.424	88.22
CLIP L14	0.429	91.90	0.441	90.62	0.442	90.16
OpenCLIP B16	0.451	90.92	0.421	89.20	0.419	88.26
OpenCLIP L14	0.429	92.20	0.443	90.80	0.440	90.04
DINOv3	0.513	84.48	0.362	81.56	0.361	80.80
Delta	171.73	0.00	0.001	0.12	0.089	8.58
5D Uniform	0.623	76.06	0.200	58.04	0.237	69.30
Bimodal Gaussian	16.380	15.88	0.074	16.36	0.068	15.16

### Hypothesis 2: Induced Normal Embeddings (INE)

Let  $Z \sim \mathcal{N}(0, I_D)$  be a random variable defined on the UNE. We hypothesize that latent representations learned by modern encoders and generative models are induced by the UNE: for each model  $i$  there exists a linear map  $C_i \in \mathbb{R}^{d_i \times D}$  and a noise term  $\epsilon_i$  such that its latent code  $\hat{Z}_i$  satisfies

$$\hat{Z}_i = C_i Z + \epsilon_i. \quad (1)$$

Consequently, each  $\hat{Z}_i$  is (approximately) Gaussian and can be viewed as a noisy linear projection of  $Z$ .

The UNE and INE hypotheses align with the Platonic Representation Hypothesis [26], but extend it in several important ways. First, they explicitly state the Gaussianity of the underlying distribution, and state the correlation between the real distribution and the distribution of observations. Second, they unify not only encoders but both families of encoders and generative models. Lastly, since INEs are noisy linear projections of the UNE, and we have access to them, we can extrapolate properties such as linear separability.

**Relation between INEs and UNE.** INEs do not achieve the ideal Gaussian latent space. However, they contain a strong normal core: many latent directions behave as nearly Gaussian, while others capture redundancy or noise. While generative models (e.g., diffusion models) are trained to sample from a Gaussian latent prior, for representation models this happens without explicit normality constraints. Foundation encoders (CLIP, OpenCLIP, DINOv3 [41, 45, 54]) empirically push embeddings toward smooth and isotropic distributions. Consequently, both representation models and generative models naturally form latent spaces where

“Gaussian-like” directions coexist with nuisance dimensions. This phenomenon is experimentally verified in Table 1, where we assess eight models: three generators from the Stable Diffusion family [38, 49] and five encoders (two CLIP variants, two OpenCLIP variants, and DINOv3). Across most models, more than 90% of latent dimensions satisfy Gaussianity according to standard normality tests (Anderson-Darling and D’Agostino-Pearson [1, 16]), confirming that learned latents already approximate the normal structure predicted by the UNE hypothesis. Experimental details are provided in Section 4.1.

### 3.2. Semantic directions

A key property of Gaussian latent spaces is that Gaussian variables interact *linearly*. If a latent code  $Z \in \mathbb{R}^d$  is standard normal and a semantic attribute  $Y \in \mathbb{R}$  is jointly Gaussian with  $Z$ , then the conditional expectation of  $Y$  given the code is linear:

$$\mathbb{E}[Y | Z] = w^\top Z + b, \quad (2)$$

for some  $w \in \mathbb{R}^d$  and  $b \in \mathbb{R}$ . This follows directly from the closed form of the multivariate Gaussian conditional distribution [12]. In this case, semantic variation corresponds to a linear direction in the latent space.

Many semantic attributes (e.g., age, height, smile intensity) behave approximately Gaussian when observed over a population: real-world measurements that arise from many small sources of variation tend to cluster around a mean and spread smoothly. In a Gaussianized latent space, such attributes align with linear directions  $w$ , making them effectively modeled by linear probes. This motivates linear classifiers or regressors in latent space, which is well established for representation models (e.g., CLIP). We further find that the same linear separability emerges in generative

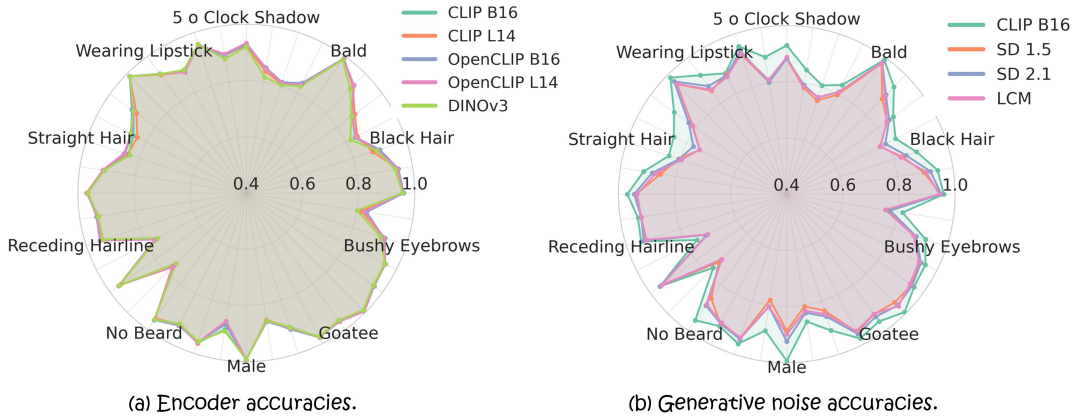


Figure 3. **Classification probing in latent spaces.** We train linear attribute classifiers (logistic regression) on latent representations from different models and evaluate accuracy on 40 CelebA attributes. (a) CLIP variants, OpenCLIP variants, and DINOv3 achieve nearly identical performance across attributes, demonstrating that semantic information is linearly accessible. (b) DDIM-inverted noise latents from SD 1.5, SD 2.1, and LCM achieve accuracy highly correlated with a strong encoder baseline (CLIP-B/16), despite originating from diffusion noise rather than semantic encoders. For clarity, only 10 representative attribute names are displayed.

latent spaces such as DDIM-inverted noise, validated empirically in Figure 3; details in Section 4.2.

**Linear editing in latent spaces.** With Gaussian latents and approximately Gaussian attributes, semantic changes often correspond to moving along linear directions. This behavior is not limited to ideal UNEs: representation and generative models whose latents only approximate Gaussianity (e.g., diffusion noise through DDIM inversion) exhibit the same effect: linear probes reveal interpretable semantic directions. In this setting, semantic editing corresponds to moving along a linear path:

$$\tilde{z} = z + \alpha w, \quad (3)$$

where  $w$  is the normal of the learned linear decision boundary and  $\alpha$  controls edit strength. We demonstrate this simple linear editing in the DDIM-inverted space in Figure 4; details in Section 4.3.

**Mitigating spurious features.** Semantic directions are not always perfectly disentangled: a direction estimated for one attribute may partially align with another, causing edits to change unintended properties. To mitigate this, we edit along an *orthogonalized* direction that removes the observed unintended changes by projecting the semantic direction into the null space of the spurious direction. Formally, let  $w_1, w_2 \in \mathbb{R}^d$  be linear directions for two attributes. Changing attribute  $B_1$  without affecting attribute  $B_2$  can be formalized as:

$$\tilde{w}_1 = w_1 - \frac{w_2 w_2^\top}{w_2^\top w_2} w_1. \quad (4)$$

An illustration and examples of this simple mitigation strategy are presented in Figure 5; see details in Section 4.3.

### 3.3. Mapping between models and shared spaces

**Mapping between models.** Each INE can be viewed conceptually as a noisy linear transformation of the same underlying normal latent space. Under this view, different models do not learn unrelated representations, they learn different linear embeddings of the same latent geometry. Therefore, moving from one model’s latent space to another should require only a linear mapping, with deviations attributable to noise or unused dimensions rather than fundamentally different structure. While prior work has separately reported linear mapping within model families (encoders and generators), our hypotheses link both within a single latent framework. This suggests a direct correspondence between generative latents (e.g., DDIM-inverted diffusion noise) and representation embeddings (e.g., encoders). We demonstrate this cross-family alignment in Table 2; experimental details are in Section 4.2.

**Recovering the shared subspace of multiple INEs.** Given  $m$  models, each produces a learned latent representation of the same  $n$  images. Although these representations differ in dimensionality and contain noise or redundant directions, they are assumed to originate from the same underlying latent structure, the UNE (see Equation (1)). Our goal is to recover a shared  $k$ -dimensional latent space that all models “agree on”.

Let  $\hat{Z}_i \in \mathbb{R}^{n \times d_i}$  denote the latent codes of model  $i$  (rows are samples, columns are centered features). We seek a shared  $k$ -dimensional representation  $X \in \mathbb{R}^{n \times k}$  such that each model can linearly explain this same latent structure via some matrix  $A_i \in \mathbb{R}^{d_i \times k}$ :

$$\hat{Z}_i A_i \approx X \quad \forall i = 1, \dots, m. \quad (5)$$

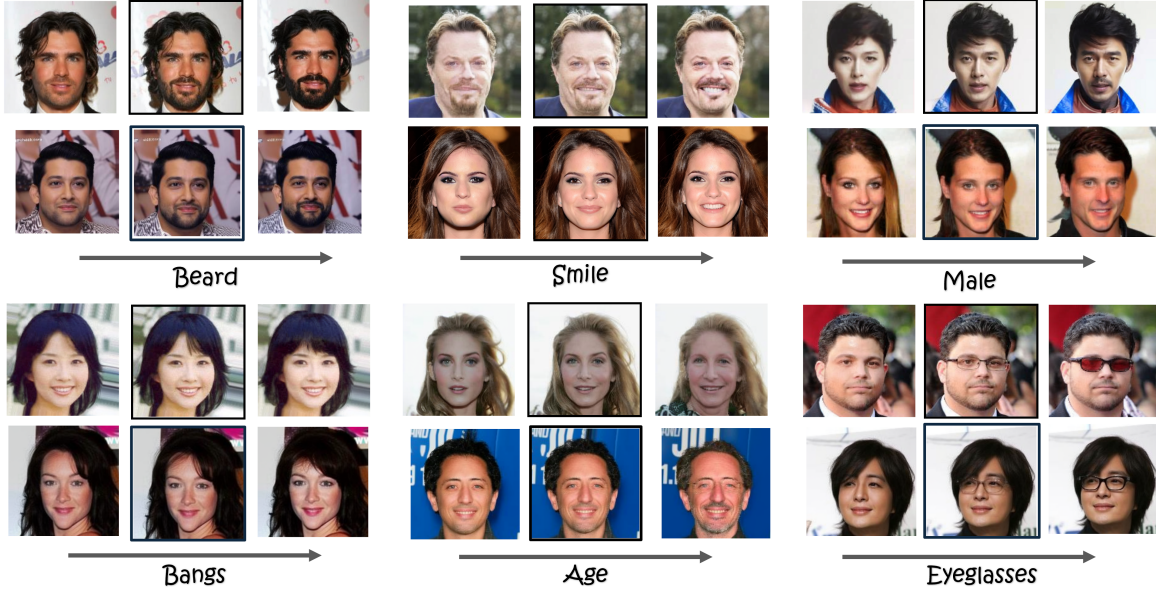


Figure 4. **Linear latent editing across semantic attributes.** For DDIM-inverted SD 1.5 latents, we move along linear classifier-derived semantic directions ( $\tilde{z} = z + \alpha w$ ). Each row shows decreasing (left) and increasing (right) attribute intensity as  $\alpha$  varies, middle image in each triplet is the original image. No prompts or model tuning are used; edits are controlled solely by linear shifts in the latent space.

Under the INE hypothesis (Equation (1)), each  $\hat{Z}_i$  is an approximately linear projection of the UNE. We therefore treat  $X$  as a  $k$ -dimensional proxy for this core space, and the  $A_i$  as approximate “inverse” projections that recover  $X$  from each INE. This leads to the following objective:

$$\begin{aligned} \min_{X, \{A_i\}_{i=1}^m} \sum_{i=1}^m \|\hat{Z}_i A_i - X\|_F^2 + \lambda_i \|A_i\|_F^2 \\ \text{s.t. } X^\top X = I, \quad \mathbf{1}^\top X = 0. \end{aligned} \quad (6)$$

The constraints  $X^\top X = I$  and  $\mathbf{1}^\top X = 0$  enforce centered features and identity covariance (up to scale) for the shared space, making  $X$  an approximate instance of the form predicted by the UNE hypothesis.  $\lambda_i > 0$  are chosen regularization parameters. This objective corresponds to the MAXVAR formulation of Generalized Canonical Correlation Analysis (GCCA) [28]. It admits a closed-form solution: the matrix  $X$  is obtained as the  $k$  eigenvectors corresponding to the smallest eigenvalues of a matrix constructed from  $\{\hat{Z}_i\}$  and  $\{\lambda_i\}$ . We note that the simplest form of GCCA sets  $\lambda_i = 0$  and drops the centering constraint  $\mathbf{1}^\top X = 0$ , but these nuances do not change the essential solution method. Our particular implementation is a hybrid approach that first optimizes  $A_i$  in Equation (6) in closed form in terms of  $X$ , and then optimizes  $X$  with  $\lambda_i = 0$  for all  $i$ .

Intuitively, this procedure identifies the *intersection* of multiple INEs. While it may not recover the full UNE, it

extracts the portion of the latent structure consistently expressed across all models, and should be viewed as an initial construction among many possible alternatives.

## 4. Experiments

We curate *NoiseZoo*, a dataset of per-image latents, and evaluate along three axes: (i) linear classification within and across latent spaces; (ii) controllable linear editing along probe-derived directions; (iii) recovery of a shared  $k$ -dimensional core via a multi-view estimator.

### 4.1. NoiseZoo construction

We use the CelebA [37] validation set ( $\sim 19k$  images, split into  $15k$  training and  $4k$  test samples). For each image, we extract latent representations from five vision encoders: CLIP ViT-L/14, CLIP ViT-B/16, OpenCLIP ViT-L/14, OpenCLIP ViT-B/16, and DINOv3 [54]. CLIP and OpenCLIP are contrastive image-text models trained on large-scale captioned datasets, whereas DINOv3 is trained purely on images using a self-supervised objective. In addition, we obtain DDIM-inverted noise latents from three generative models in the Stable Diffusion family: SD 1.5, SD 2.1, and LCMv7 [38, 49]. SD 1.5 and SD 2.1 differ in training data and text encoders, while LCMv7 is trained under the Latent Consistency Model objective, which enables few-step sampling and induces a different geometry in the noise latent space. Across models, encoder latents are moderately sized ( $500\text{--}1k$  dimensions), whereas

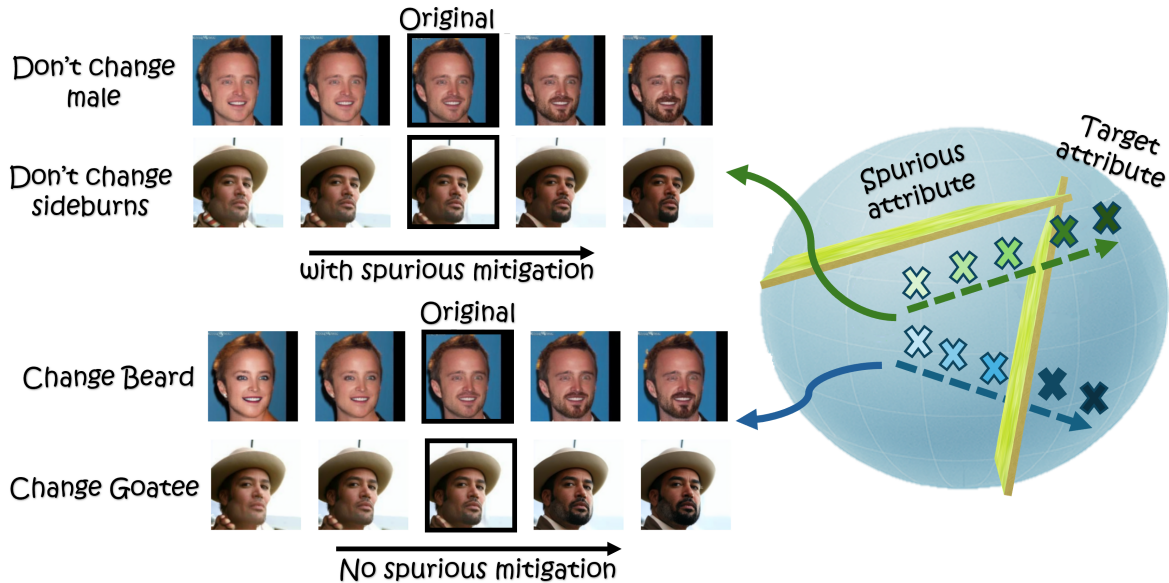


Figure 5. **Removing spurious attribute correlations.** Edits performed using the raw semantic direction (bottom) unintentionally modify a correlated attribute (e.g., adding a goatee also changes facial structure). Using the orthogonalized direction (top) from Equation (4) isolates the target attribute while suppressing the spurious one, yielding clean, disentangled edits.

Table 2. **Transferred latent evaluations.** For each generative model, we linearly transfer its latents into encoder latent spaces and evaluate (i) geometric similarity (MSE, cosine similarity) and (ii) downstream attribute prediction accuracy using fixed encoder-trained classifiers. An insignificant drop in accuracy (less than 0.3%) and high similarity (high cosine similarity, low MSE) confirm that generative noise latents can be linearly aligned to encoder spaces while preserving predictive structure.

Model	MSE ↓			Cosine Similarity ↑			Accuracy Drop (pp) ↓		
	CLIP B16	OpenCLIP L14	DINOv3	CLIP B16	OpenCLIP L14	DINOv3	CLIP B16	OpenCLIP L14	DINOv3
SD 1.5	0.072	0.215	0.143	0.8	0.73	0.55	0.20	0.23	0.23
SD 2.1	0.073	0.219	0.146	0.80	0.72	0.54	0.14	0.19	0.20
LCM	0.07	0.21	0.137	0.81	0.74	0.57	0.00	0.00	0.00

DDIM-inverted diffusion latents have much higher dimensionality ( $\sim 16k$ ). Together, these models provide diverse generative and representation embeddings for the same underlying images. This yields *NoiseZoo*: a set of latents for every image. Details and examples are in Supp. Section A.

**Assessing Gaussianity.** We evaluate Gaussianity using Anderson-Darling, D’Agostino-Pearson, and Shapiro-Wilk tests on random 1D projections of the latent space [1, 16, 51]. For each model, we sample 250 data points, compute 5,000 random projections, and report: (i) the average test statistic; and (ii) the fraction of projections that do *not* reject normality. As shown in Table 1, generative models approach the theoretical 95% acceptance rate of Gaussian samples and encoder representations score slightly lower but remain high. We additionally examine non-Gaussian reference distributions (delta distributions, low-dimensional uniform distributions, and bimodal Gaussians) as controls. These references perform substantially worse than both

generative model and encoder representation spaces.

## 4.2. Classification in latent spaces

For each model, we train logistic-regression classifiers (training details in Supp. Section A) for the 40 CelebA attributes using its training latents, and evaluate them on the corresponding test latents. Attribute-wise accuracies for encoders and generative models are shown in Figure 3, with CLIP ViT-B/16 overlaid on the generative panel for reference. Overall, DDIM-inverted noise latents yield attribute separability only slightly below that of leading encoders, with highly correlated per-attribute behavior across models.

**Cross-space transfer.** To evaluate transferability, we learn ridge-regularized linear maps (Supp. Section A) from each generative latent space into three encoder spaces using the training split, and apply the *fixed* encoder-trained classifiers to the mapped latents of the test set. Post-transfer performance is nearly unchanged, indicating that linear align-




Figure 6. **Classification accuracy in shared latent spaces.**  $X_1$ – $X_4$  denote shared spaces computed from four latent sources, and  $X_5$  from six, with combinations detailed in Section 4.4. (a) Attribute classification accuracy as a function of latent dimension for shared spaces  $X_1$ – $X_5$ , showing strong performance even at low dimensions. (b) Linear-probe accuracy at 16 dimensions in the PCA-reduced latent spaces of each model and in the low-dimensional shared spaces  $X_1$ – $X_5$ , indicating that the shared space intersections retain comparable attribute information. (c) Retrieval-based analysis: pairwise correlations (Spearman rank correlation) of similarity vectors (computed over 10k latents) between the shared spaces  $X_1$ – $X_5$ , demonstrating that they encode highly similar underlying structure.

ment suffices for downstream prediction. Mean square error (MSE), cosine similarity, and accuracy drops are reported in Table 2, showing low error, high similarity, and minimal degradation. Additional results appear in Supp. Section B.

### 4.3. Linear editing

Using the semantic directions from our linear classifiers, we edit DDIM-inverted latents and decode the modified samples. Figure 4 shows edits across six CelebA attributes (SD 1.5), where varying intensity ( $\alpha$  in Equation (3)) smoothly increases or decreases attribute strength. Edits are local, controllable, and require no prompts or fine-tuning. To mitigate attribute entanglement, we use the orthogonalization in Equation (4); Figure 5 demonstrates that it isolates the target attribute and suppresses spurious changes. We apply this procedure across all generative models and to CLIP ViT-L/14 (usable for synthesis guidance via UnCLIP [46]). Quantitative and additional qualitative results are provided in Supp. Section B.

### 4.4. Shared latent spaces

We compute shared latent spaces  $X_i$  using the multi-view intersection method in Section 3.3. We consider four shared spaces constructed from four sources each and one from six sources; in every case, each shared space combines an equal number of encoder and generative model latents (see exact splits in Supp. Section A). Figure 6a shows that all shared spaces achieve strong attribute classification in medium to high dimensions (32–512) and degrade in a similar manner when the dimension is reduced below 16. We additionally apply PCA to each latent space individually; Figure 6b shows that, at 16 dimensions, these PCA-reduced spaces reach similar classification performance to the shared spaces. Since each shared space is the intersection of its sources, it cannot contain more information

than any single latent space. This suggests that attribute information concentrates in a small set of shared directions. To further test similarity between shared spaces, we perform a retrieval analysis on a subset of 10k images: for each test latent, we measure its cosine similarity to all other latents and obtain a similarity vector. Figure 6c reports Spearman rank correlations between these vectors across shared spaces, which are consistently high, indicating similar neighborhood structure. Overall, these results provide a preliminary *proof of concept* for the UNE hypothesis, suggesting that latent representations from both encoders and generative models retain highly similar underlying information.

## 5. Conclusions

We introduced the *Universal Normal Embedding* hypothesis, proposing that generative and representation models approximate a shared Gaussian latent geometry where semantic factors correspond to linear directions. An immediate consequence of the UNE is that generators hold linearly separable semantics, in a similar manner to encoders. Empirically, we demonstrated that DDIM-inverted noise codes and representation embeddings encode comparable semantic structure: attributes are linearly decodable in both, their probe predictions strongly agree, and noise-space classifier directions enable controllable edits without retraining or architectural changes. Current work primarily bridges the hypothesis with empirical findings. We plan to characterize the mechanisms that drive models toward UNE-like geometry in future studies. This work is a step toward a unification of representation learning and generative modeling, suggesting a shared geometric framework. We believe this viewpoint can guide both theoretical advances and the design of principled, interpretable generative systems.

## Acknowledgments

We would like to acknowledge support by the Israel Science Foundation (Grant 1472/23) and by the Ministry of Innovation, Science and Technology (Grant 8801/25).

## References

- [1] Theodore W Anderson and Donald A Darling. A test of goodness of fit. *Journal of the American statistical association*, 49(268):765–769, 1954. 4, 7
- [2] Andrea Asperti and Valerio Tonelli. Comparing the latent space of generative models. *Neural Computing and Applications*, 35(4):3155–3172, 2023. 1, 2
- [3] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18208–18218, 2022. 1
- [4] Charumathi Badrinath, Usha Bhalla, Alex Oesterling, Suraj Srinivas, and Himabindu Lakkaraju. All roads lead to Rome? exploring representational similarities between latent spaces of generative image models. *arXiv preprint arXiv:2407.13449*, 2024. 1, 2
- [5] Randall Balestriero and Yann LeCun. LeJEPa: Provable and scalable self-supervised learning without the heuristics. *arXiv preprint arXiv:2511.08544*, 2025. 2, 3
- [6] Randall Balestriero, Nicolas Ballas, Mike Rabbat, and Yann LeCun. Gaussian embeddings: How JEPAs secretly learn your data density. *arXiv preprint arXiv:2510.05949*, 2025. 3
- [7] Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Revisiting model stitching to compare neural representations. *Advances in neural information processing systems*, 34:225–236, 2021. 1, 2
- [8] Adrien Bardes, Jean Ponce, and Yann LeCun. VI-CReg: Variance-invariance-covariance regularization for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021. 3
- [9] Roy Betser, Meir Yossef Levi, and Guy Gilboa. Whitened CLIP as a likelihood surrogate of images and captions. In *Proceedings of the 42nd International Conference on Machine Learning*, Vancouver, Canada, 2025. PMLR. 3
- [10] Roy Betser, Eyal Gofer, Meir Yossef Levi, and Guy Gilboa. Infonce induces gaussian distribution. In *International Conference on Learning Representations (ICLR)*, 2026. 2, 3
- [11] Roy Betser, Omer Hofman, Roman Vainshtein, and Guy Gilboa. General and domain-specific zero-shot detection of generated images via conditional likelihood. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7809–7820, 2026. 2
- [12] Christopher M Bishop and Nasser M Nasrabadi. *Pattern recognition and machine learning*. Springer, 2006. 2, 4
- [13] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 1, 2
- [14] Siyi Chen, Huijie Zhang, Minzhe Guo, Yifu Lu, Peng Wang, and Qing Qu. Exploring low-dimensional subspace in diffusion models for controllable image editing. *Advances in neural information processing systems*, 37:27340–27371, 2024. 3
- [15] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PmLR, 2020. 1
- [16] Ralph D’Agostino and Egon S Pearson. Tests for departure from normality. Empirical results for the distributions of  $b_2$  and  $\sqrt{b_1}$ . *Biometrika*, 60(3):613–622, 1973. 4, 7
- [17] Yusuf Dalva and Pinar Yanardag. NoiseCLR: A contrastive learning approach for unsupervised discovery of interpretable directions in diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 24209–24218, 2024. 3
- [18] Imant Daunhawer, Alice Bizeul, Emanuele Palumbo, Alexander Marx, and Julia E Vogt. Identifiability results for multimodal contrastive learning. *arXiv preprint arXiv:2303.09166*, 2023. 2, 3
- [19] Aleksandr Ermolov, Aliaksandr Siarohin, Enver Sangineto, and Nicu Sebe. Whitening for self-supervised representation learning. In *International conference on machine learning*, pages 3015–3024. PMLR, 2021. 3
- [20] Ian J 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. 1, 2
- [21] René Haas, Inbar Huberman-Spiegelglas, Rotem Mulayoff, Stella Graßhof, Sami S Brandt, and Tomer Michaeli. Discovering interpretable directions in the semantic latent space of diffusion models. In *2024 IEEE 18th International Conference on Automatic Face and Gesture Recognition (FG)*, pages 1–9. IEEE, 2024. 3
- [22] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. GANSpace: Discovering interpretable GAN controls. *Advances in neural information processing systems*, 33:9841–9850, 2020. 3
- [23] Omer Ben Hayun, Roy Betser, Meir Yossef Levi, Levi Kassel, and Guy Gilboa. Training-free detection of generated videos via spatial-temporal likelihoods. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026. 2
- [24] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 3
- [25] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 1, 2
- [26] Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation hypothesis. *arXiv preprint arXiv:2405.07987*, 2024. 1, 2, 3, 4
- [27] Rishi Jha, Collin Zhang, Vitaly Shmatikov, and John X Morris. Harnessing the universal geometry of embeddings. *arXiv preprint arXiv:2505.12540*, 2025. 1, 2, 3

- [28] Jon R Kettenring. Canonical analysis of several sets of variables. *Biometrika*, 58(3):433–451, 1971. 6
- [29] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1, 2
- [30] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. Similarity of neural network representations revisited. In *International conference on machine learning*, pages 3519–3529. PMIR, 2019. 1, 2
- [31] Mingi Kwon, Jaeseok Jeong, and Youngjung Uh. Diffusion models already have a semantic latent space. In *The Eleventh International Conference on Learning Representations*, 2023. 3
- [32] Zorah Löhner and Michael Moeller. On the direct alignment of latent spaces. In *Proceedings of UniReps: the First Workshop on Unifying Representations in Neural Models*, pages 158–169. PMLR, 2024. 1, 2
- [33] Meir Yossef Levi and Guy Gilboa. The double ellipsoid geometry of CLIP. In *Proceedings of the 42nd International Conference on Machine Learning*, Vancouver, Canada, 2025. PMLR. 3
- [34] Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. On the sentence embeddings from pre-trained language models. *arXiv preprint arXiv:2011.05864*, 2020. 2
- [35] Victor Weixin Liang, Yuhui Zhang, Yongchan Kwon, Serena Yeung, and James Y Zou. Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. *Advances in Neural Information Processing Systems*, 35:17612–17625, 2022. 3
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. 1, 2
- [37] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 6
- [38] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023. 4, 6
- [39] Valentino Maiorca, Luca Moschella, Antonio Norelli, Marco Fumero, Francesco Locatello, and Emanuele Rodolà. Latent space translation via semantic alignment. In *Advances in Neural Information Processing Systems*, pages 55394–55414. Curran Associates, Inc., 2023. 1, 2
- [40] Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. Linearly mapping from image to text space. *arXiv preprint arXiv:2209.15162*, 2022. 1, 2
- [41] mlfoundations. OpenCLIP. [https://github.com/mlfoundations/open\\_clip](https://github.com/mlfoundations/open_clip), 2021. OpenCLIP: Open reproduction of CLIP training, Github page. 4
- [42] Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6038–6047, 2023. 3
- [43] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020. 3
- [44] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. StyleCLIP: Text-driven manipulation of StyleGAN imagery. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 2085–2094, 2021. 1
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021. 1, 2, 4
- [46] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022. 8
- [47] Patrik Reizinger, Alice Bizeul, Attila Juhos, Julia E Vogt, Randall Balestriero, Wieland Brendel, and David Klindt. Cross-entropy is all you need to invert the data generating process. *arXiv preprint arXiv:2410.21869*, 2024. 2, 3
- [48] Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR, 2015. 2
- [49] 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. 1, 4, 6
- [50] Simon Schrodi, David T Hoffmann, Max Argus, Volker Fischer, and Thomas Brox. Two effects, one trigger: On the modality gap, object bias, and information imbalance in contrastive vision-language representation learning. *arXiv preprint arXiv:2404.07983*, 2024. 3
- [51] Samuel Sanford Shapiro and Martin B Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3-4):591–611, 1965. 7
- [52] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of GANs for semantic face editing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9243–9252, 2020. 3
- [53] Peiyang Shi, Michael C Welle, Márten Björkman, and Danica Kragic. Towards understanding the modality gap in clip. In *ICLR 2023 workshop on multimodal representation learning: perks and pitfalls*, 2023. 3
- [54] Oriane Siméoni, Huy V. Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, Francisco Massa, Daniel Haziza, Luca Wehrstedt, Jianyuan Wang, Timothée Darcet, Théo Moutakanni, Leonel Sentana, Claire Roberts, Andrea Vedaldi, Jamie Tolan, John Brandt, Camille Couprie, Julien Mairal, Hervé Jégou, Patrick Labatut, and Piotr Bojanowski. DINOv3, 2025. 4, 6
- [55] Bowen Song, Zecheng Zhang, Zhaoxu Luo, Jason Hu, Wei Yuan, Jing Jia, Zhengxu Tang, Guanyang Wang, and Liyue

- Shen. Ccs: Controllable and constrained sampling with diffusion models via initial noise perturbation. *arXiv preprint arXiv:2502.04670*, 2025. [3](#)
- [56] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. [2](#)
- [57] Kirill Tyshchuk, Polina Karpikova, Andrew Spiridonov, Anastasiia Prutianova, Anton Razzhigaev, and Alexander Panchenko. On isotropy of multimodal embeddings. *Information*, 14(7):392, 2023. [2](#)
- [58] Li Wang, Boyan Gao, Yanran Li, Zhao Wang, Xiaosong Yang, David A Clifton, and Jun Xiao. Exploring the latent space of diffusion models directly through singular value decomposition. *arXiv preprint arXiv:2502.02225*, 2025. [3](#)
- [59] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. [1](#), [2](#)
- [60] Tongzhou Wang and Phillip Isola. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*, pages 9929–9939. PMLR, 2020. [3](#)
- [61] Can Yaras, Siyi Chen, Peng Wang, and Qing Qu. Explaining and mitigating the modality gap in contrastive multimodal learning. *arXiv preprint arXiv:2412.07909*, 2024. [3](#)
- [62] Jure Zbontar, Li Jing, Ishan Misra, Yann LeCun, and Stéphane Deny. Barlow twins: Self-supervised learning via redundancy reduction. In *International conference on machine learning*, pages 12310–12320. PMLR, 2021. [3](#)
- [63] Huijie Zhang, Jinfan Zhou, Yifu Lu, Minzhe Guo, Liyue Shen, and Qing Qu. The emergence of reproducibility and consistency in diffusion models. In *International Conference on Machine Learning*. PMLR, 2024. [1](#), [2](#)
- [64] Roland S Zimmermann, Yash Sharma, Steffen Schneider, Matthias Bethge, and Wieland Brendel. Contrastive learning inverts the data generating process. In *International conference on machine learning*, pages 12979–12990. PMLR, 2021. [1](#), [2](#), [3](#)