

# Latent Space Imaging

Matheus Souza

Yidan Zheng

Kaizhang Kang

Yogeshwar Nath Mishra

Qiang Fu

Wolfgang Heidrich

KAUST

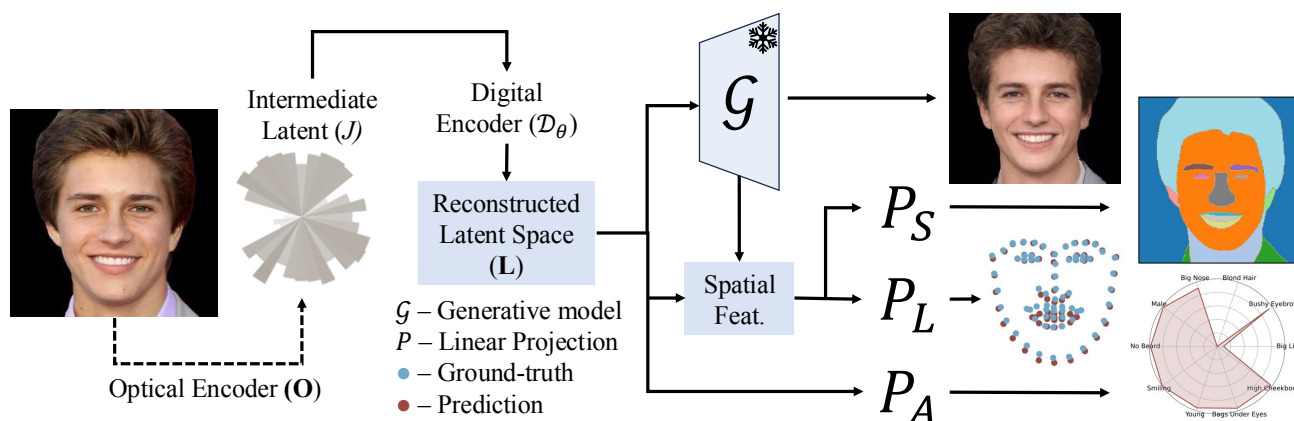


Figure 1. We propose an extremely-compressed imaging paradigm called Latent Space Imaging (LSI). The optical encoder (O) projects the real signal into a compressed set of measurements. A digital encoder ( $\mathcal{D}_\theta$ ) then maps this signal to the latent space (L) of a frozen generative model (G), enabling image reconstruction. The L can also be linearly projected (P) to perform downstream tasks directly—such as facial segmentation ( $P_S$ ), landmark detection ( $P_L$ ), and attribute classification ( $P_A$ )—without requiring image reconstruction or a complex new model.

## Abstract

Digital imaging systems have traditionally relied on brute-force measurement and processing of pixels arranged on regular grids. In contrast, the human visual system performs significant data reduction from the large number of photoreceptors to the optic nerve, effectively encoding visual information into a low-bandwidth latent space representation optimized for brain processing. Inspired by this, we propose a similar approach to advance artificial vision systems.

*Latent Space Imaging introduces a new paradigm that combines optics and software to encode image information directly into the semantically rich latent space of a generative model. This approach substantially reduces bandwidth and memory demands during image capture and enables a range of downstream tasks focused on the latent space.*

We validate this principle through an initial hardware prototype based on a single-pixel camera. By implementing an amplitude modulation scheme that encodes into the gen-

erative model’s latent space, we achieve compression ratios ranging from 1:100 to 1:1000 during imaging, and up to 1:16384 for downstream applications. This approach leverages the model’s intrinsic linear boundaries, demonstrating the potential of latent space imaging for highly efficient imaging hardware, adaptable future applications in high-speed imaging, and task-specific cameras with significantly reduced hardware complexity.

## 1. Introduction

Conventional imaging systems are based on brute-force sampling of scene information on regular grids. In contrast, biological vision systems operate very differently: for example the human visual system has around 120 million rods and 7 million cones [40], but the information from these photoreceptors is distilled by retinal processing into neural signals that are transmitted over the optic nerve with only 0.7-1.7 million axons [23]. Therefore, our vision employs compression to encode image information into a la-

tent space compatible with the neural architecture of the human brain.

Here we propose Latent Space Imaging (LSI), a new paradigm that aims to bring this concept to the realm of cameras and artificial vision systems. Specifically, we propose to utilize the combination of linear optical encoding for drastic data reduction, paired with minimal non-linear computational processing to directly encode image information into the latent space of a generative model. This approach extremely reduces the amount of sensor data required for downstream tasks such as face segmentation, facial landmark detection, or image reconstruction.

To showcase this framework, we focus here on the rich latent space of a generative model in the form of a StyleGANXL [26–28, 49] trained on human faces. StyleGAN models are renowned for their ability to synthesize highly realistic face images, while supporting editing and image-to-image translation applications. Their semantically rich, disentangled, well-structured, and *compact* latent space makes them an ideal architecture for LSI. This structured representation facilitates key downstream applications, including landmark detection, facial feature segmentation, and the extraction of facial attributes such as age, gender, smile/no smile, and beard/no beard. Furthermore, we demonstrate its effectiveness in the particularly challenging task of reconstructing full human face images from a very small number of samples.

While there are a number of possible ways to optically implement the LSI paradigm, our first prototype utilizes the single-pixel imaging (SPI) concept [13, 29, 55, 71]. A discussion of alternative implementations [6, 37, 63, 64] can be found in the Supplement.

In summary, we make the following contributions:

- the introduction of Latent Space Imaging as a new paradigm for very low bandwidth image capture;
- a proof-of-concept hardware prototype for LSI utilizing the SPI framework, featuring a programmable optical system for easy and extensive experimentation,
- an initial latent space for human faces based on a StyleGAN’s family of generative model; and
- the demonstration of a range of downstream applications on this latent space with real hardware experimentation.

We also refer the readers to the Supplement for a detailed discussion of hardware and software choices for the LSI paradigm. Our implementation is available at <https://github.com/vccimaging/latent-imaging>.

## 2. Related Work

### 2.1. Optics and Algorithm Co-design

Latent Space Imaging builds and improves upon the recent trend for co-designing both optics and reconstruction algorithms. This joint optimization has yielded favorable

outcomes in computational imaging and downstream computer vision tasks. These include color image restoration [4], microscopy [33, 42, 51], monocular depth imaging [5, 16, 17, 65], super-resolution, extended depth of field [53, 57], time-of-flight imaging [8], and high-dynamic range imaging [56], among others. The enhancements observed are attributed to the ability of such systems to *optically encode* relevant scene information into sensor data, with the computational post-processing acting as a *decoder*. The crucial innovation in LSI is that the code represents the semantically expressive yet efficiently encoded latent space of a generative model, with that model acting as the decoder module.

### 2.2. Compressive Sensing and Minimal Cameras

Our approach aligns with the principles of Compressive Sensing (CS) [12], which harnesses sparsity and structural priors to recover signals from minimal samples. Traditional CS methods emphasize optimization-based reconstruction, while recent advancements integrate deep learning, combining sparsity-driven priors with neural network capabilities. This first group of methods has led to diverse architectures, including CNNs, transformers, and deep unfolding models [41, 45, 54, 60–62, 69], bridging the gap between classical CS and data-driven approaches. Further, generative priors, forming a second group, offer a robust alternative by leveraging latent spaces to enhance image quality and structure [3, 22, 24, 31, 32], showing the potential of learned priors to capture dependencies within compressed data. A third group advances these concepts in real-world applications through single-pixel imaging [13, 19, 20, 29, 48, 55, 68, 71], using quantized patterns, binary random masks, Fourier domain methods, and co-optimized mask-reconstruction pipelines. Lastly, a fourth group [35, 44] presents minimalist cameras with optimized freeform pixels for task-specific applications like workspace monitoring.

*Latent Space Imaging* stands apart from both the first and second groups by emphasizing realistic physical deployment, requiring attention to numerical constraints, quantization, light efficiency, and similar practical elements. Unlike the second group, which often employs iterative optimization, our approach reconstructs images in a single forward pass. Additionally, we distinguish our work from these groups by targeting latent space reconstruction itself. Our framework uniquely leverages once-optimized masks for multi-model purposes, avoiding the conventional use of latent space solely as an optimization prior and without optimizing the masks/pixels for multiple tasks. Finally, we set ourselves apart from the fourth group design by utilizing a latent space to enable fine-grained tasks, such as face identification, with a smaller number of *freeform* samples than traditional cameras.

### 3. Latent Space Imaging

Traditional imaging captures a regular  $m \times n$  grid of measurements, while compression strategies typically use a collection of random linear codes to leverage the information redundancy inherent in real images in combination with a reconstruction software module. In contrast, we demonstrate that meaningful information can be significantly more compressed in the latent space of generative models while providing sufficient information for downstream processing tasks.

As depicted by Fig. 1, our objective is to learn an optically encoded mapping

$$\mathbf{L} = \mathcal{D}_\theta(\mathbf{O} \cdot I), \quad (1)$$

where the input image  $I \in \mathbb{R}^{mn}$  is mapped to a tensor  $\mathbf{L} \in \mathbb{R}^{k \times l}$  representing the latent space of a generative model. This mapping is achieved through the combination of a co-designed (linear) optical encoder  $\mathbf{O} \in \mathbb{R}^{d \times mn}$  and a digital (nonlinear) encoder  $\mathcal{D}_\theta$ , with  $\theta$  representing the set of parameters that configure the digital encoder. Where  $d$  is the dimension of the compressed measurement,  $l$  is the number levels from coarse to fine inside the generative models as well as in the latent space, and  $k$  represents latent dimension.

The optical encoder linearly projects the image to an intermediate vector  $J \in \mathbb{R}^d$  ( $\mathbf{O} \cdot I$ ) representing the actual measurements of the signal. This substantial compression ratio is achieved by optimizing the projection alongside other components during the image reconstruction process. Instead of merely matching pixel values, the process targets the generative model's latent representation  $\mathbf{L}$  of the original signal. By focusing on this compressed space, the essential features of the data distribution within a particular domain are captured, effectively eliminating redundancies inherent in the data.

However, due to the complex non-linear operations intrinsic to generative models and the impracticality of realizing these non-linearities optically, an additional digital encoder  $\mathcal{D}_\theta$  is necessary to align  $J$  with the entire latent space of interest. This alignment is achieved through a deep neural network (DNN), that enhances the representation by effectively addressing the non-linearities. The digital encoder  $\mathcal{D}_\theta$  expands  $J$  to  $\mathbf{L} \in \mathbb{R}^{512 \times 18}$  (see Fig. 1), aligning it with the dimensions of the chosen generative model. In our work, StyleGANXL [49] is employed since it performs robustly within specific domains, and its mechanisms for inversion to latent space are well-studied [66].

Specifically, the digital encoder employs a multi-level network inspired by advancements in the inversion domain [47]. This approach leverages the detailed structural levels of StyleGANXL. In this framework, the optically encoded vector  $J$  passes through a series of linear layers and

attention mechanisms, each followed by activation functions. The depth of these stacks varies to match the required level of detail. For the pre-trained StyleGANXL on human faces, there are 18 levels corresponding to different resolutions. The number of stacks increases progressively from coarse to fine, aligning with low to high resolutions, respectively.

After acquiring all detail levels, we obtain  $\mathbf{L} \in \mathbb{R}^{512 \times 18}$ . A mixer block, inspired by [38], is then used to learn a weighted mixture of information across these levels, which is subsequently fed into the generative model  $\mathcal{G}$  to output the image  $I'$

$$I' = \mathcal{G}(\mathbf{L}), \quad (2)$$

where  $I' \in \mathbb{R}^{mn \times 3}$ , with  $m$  and  $n$  determined by the capabilities of the generative model  $\mathcal{G}$ . The value 3 corresponds to the RGB color channels which are discussed in further detail in the Supplement.

A pivotal aspect of the LSI framework is the identification of the optimal latent space to serve as ground-truth during the optimization process. To accomplish this, we trained an encoder,  $\mathcal{E}$ , following the method proposed by [47] for accessing the StyleGANXL latent space. This encoder provides an approximate ground-truth latent representation,  $\mathcal{E}(I)$ , which serves as the target for optimizing each entry of the  $\mathbf{O}$  matrix and the  $\mathcal{D}_\theta$ . The optimization is driven by the loss function  $\mathcal{L}_{lat}$

$$\mathcal{L}_{lat} = \|\mathcal{D}_\theta(\mathbf{O} \cdot I) - \mathcal{E}(I)\|_1. \quad (3)$$

Making the encoding  $\mathbf{L}$  close to  $\mathcal{E}(I)$  guarantees that it is inside the generative model average space and diminishes hallucination caused by only matching an  $\ell_2$  norm at the pixel level. To further enhance facial fidelity, we incorporate additional terms in the loss function. Detailed information on these loss functions can be found in the Supplement.

### 4. Linear Learning for Downstream Tasks

In the LSI pipeline, we use a hybrid encoding setting to generate  $\mathbf{L}$ , enabling high-level vision tasks with minimal training and a very limited number of measurements. Previous research has examined the semantics of GAN models, uncovering latent space linear boundaries that distinguish different image attributes [1, 9, 52, 67]. In particular, [67] demonstrated that linear transformations can efficiently extract semantic information from feature maps without the need for image reconstruction or specialized encoders for various tasks.

Our approach involves performing three different high-level facial analysis tasks using simple linear transformations *from a single latent space*  $\mathbf{L}$ . Unlike prior studies, we push compression limits by assessing the minimum measurements needed to derive meaningful facial semantics

through linear projections  $P$  from the same latent space. Once  $\mathbf{O}$  and  $\mathcal{D}_\theta$  are optimized, they are kept frozen and different  $P$  are optimized for downstream applications such as:

**Attribute Classification** predicts facial attributes (*e.g.*, age, beard, smile, gender) using a linear feature mapper  $P_A$ , which is implemented as a fully connected layer that projects  $\mathbf{L} \in \mathbb{R}^{512 \times 18}$  to  $\mathbb{R}^{40}$ , representing 40 distinct facial characteristics.

**Face Segmentation** provides pixel-wise segmentation across facial regions (*e.g.*, nose, eyes, face, mouth). To accurately capture spatial structure, we employ multi-level feature extraction from the generative model, combining these features through a linear projection block,  $P_S$ , which utilizes bilinear interpolation alongside convolutional layers. The generative model  $\mathcal{G}$  offers feature maps at various resolutions, spanning from coarse ( $36 \times 36$ ) to fine ( $256 \times 256$ ). We extract features across six distinct levels, representing a comprehensive range of resolutions. This multi-level feature integration allows for precise semantic pixel assignment, effectively addressing the spatial limitations of raw  $\mathbf{L}$  features, which alone lack the fine-grained detail essential for high-accuracy segmentation.

**Landmark Detection** yields the coordinates of key facial points within an image. Unlike segmentation, where multiple features are extracted, we focus here on a single coarse feature map from the generative model to capture spatial structures. This coarse feature is denoted by  $\in \mathbb{R}^{1024 \times 36 \times 36}$  and is modulated by the latent space,  $\mathbf{L}$  itself. A convolutional layer, combined with bilinear interpolation, projects and downsamples this feature to  $\mathbb{R}^{68 \times 8 \times 8}$ . Finally, a linear layer reduces the spatial dimensions, representing the 68 key points as 2D coordinates.

## 5. Hardware Prototype

Realizing the LSI framework (as detailed in Fig. 2) involves implementing the optical encoder hardware along with the digital encoder software module. Among a range of different alternatives to realize this linear layer in real-world, we opt for the single-pixel camera approach as the hardware framework for our prototype, since it allows for programmable adjustments of the compression ratio for experimentation. The Supplement discusses alternative options for physical realization of the LSI concept.

A schematic of our setup is depicted in Fig. 2. In this system, a scene is projected onto a Digital Micromirror Device (DMD TI DLP4500) using an objective lens. The image is then spatially modulated by the optimized patterns. The modulated output is captured by a relay lens, which focuses it onto a single-pixel detector (ThorLabs PDA100A2). This detector acts as an integrator, measuring a single value that represents the integrated intensity of the entire modulated image. A computer controls and synchronizes the op-

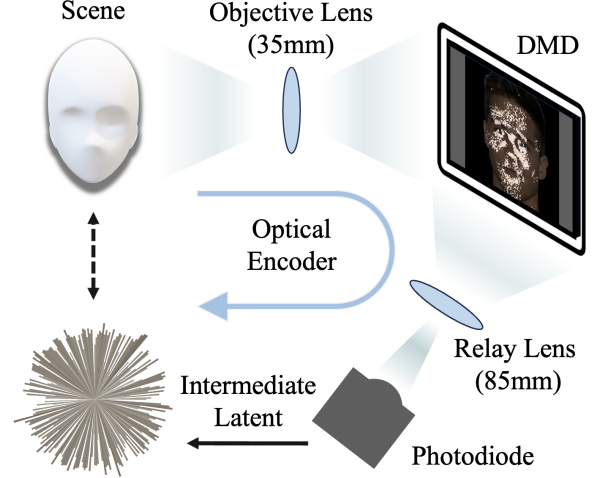


Figure 2. We illustrate one possible implementation of the Latent Space Imaging technique using a single-pixel framework. An objective lens focuses the image onto the Digital Micromirror Device (DMD), which is responsible for implementing the learned mask to spatially modulate the incoming signal. This is followed by a relay lens, which focuses the modulated signal onto a photodiode (SPD) responsible for integrating the signal. Utilizing time multiplexing, we can retrieve the necessary measurements.

eration of the DMD and the measurement signals, enabling the time-multiplexing of all measurements that compose the compressed signal.

In the single-pixel camera implementation of the LSI framework,  $\mathbf{O}$  spatially modulates the incoming image  $I$  through matrix multiplication (denoted by  $\cdot$  in Eq. 1). Each row of the optical encoder corresponds to one learned mask, and each column represents one measurement. Each mask corresponds to a single pixel, where its forms are learned for a specific domain/imaging purpose. Having the hardware in the loop during training involves maintaining  $\mathbf{O}$ 's entries as positive, bounded between  $[0, 1]$ , and quantized. In order to maintain the proposed setup as minimal as possible, we utilize binary optical masks - corresponding to quantizing its values to 0 or 1. Such decision enables reaching high frame-rate on light modulators like Digital Mirror Devices (DMDs).

With this framework established, we further improve the system's light efficiency and sensitivity to subtle variations by enforcing a 1% intensity difference between each optical mask. This distinction is achieved by introducing an energy efficiency loss in the later stages of training, fine-tuning the encoding process to ensure each pattern is uniquely identifiable, thereby significantly enhancing the system's ability to effectively discriminate between different patterns. Fig. 3 illustrates that the masks result in patterns highlighting the system's focus on facial landmarks for specific domain encoding. Please refer to the Supplement for details.






Figure 3. Displays optimized pixel forms for 1:1024 compression ratio.

## 6. Training Details

**Optical Encoder Optimization** In optimizing  $\mathbf{O}$ , we use the straight-through estimator (STE) [21], as in VQ-VAEs [39, 46, 58]. Forward passes quantize and constrain  $\mathbf{O}$  values, while back-propagation allows gradients to flow through the non-quantized version, facilitating effective weight updates.

**Datasets** We utilize FFHQ [26] and CelebAHQ [25], reserving 2000 images from CelebAHQ for testing. Only CelebAHQ is used for downstream applications, where Segment Anything (SAM) [34] focuses on faces, disregarding background content. The Supplement includes additional evaluations on the AFHQ dataset [7].

**Loss Functions** Alongside the latent 3 and energy 5 losses, we employ the identity loss ( $\mathcal{L}_{id}$ ), calculating cosine distance between feature maps from ArcFace [11]. We also apply the  $\ell_2$  norm ( $\mathcal{L}_{l2}$ ) and DINO/LPIPs features as perceptual loss ( $\mathcal{L}_p$ ) [70, 73].

The Supplement provides full details on hyperparameters, optimizers, loss functions, as well as downstream application losses and additional training technicalities.

## 7. Simulation Results

In our experiments we evaluate the LSI pipeline with simulated results as well as real-world measurements. We focus on the face domain understanding how the minimal imaging performs on essential facial tasks like reconstruction, identity recognition, attribute classification, segmentation, and landmark detection, analyzing the boundaries for each one of such tasks.

For quantitative results, we select two metrics for face identification: VGGFace [43] and Dlib [30]. In both cases, pre-trained face recognition models are used to extract features from the original and reconstructed images. Cosine similarity is then computed to verify identity; a similarity score below the thresholds of 0.68 and 0.07, respectively, indicates a match. These thresholds are sourced from the DeepFace [50] library. Additionally, FID [18] is also computed to compare the distribution of features between the reconstructed and the original set. Traditional metrics, such as PSNR and SSIM, are useful for evaluating pixel-level accuracy but often fall short in fully capturing human visual perception [2, 36, 70]. Further assessment and an in-depth

discussions of these metrics are provided in the Supplement.

### 7.1. Image Reconstruction

Full image reconstruction is the most demanding task for any reduced sampling scheme, especially in the case of human faces, where the human visual system is extremely sensitive to minute differences. Fig. 4 demonstrates the robust identity preservation capabilities of our model when reconstructing images at a resolution of  $256 \times 256$  across a range of measurement levels. Latent Space Imaging effectively captures essential features such as eye color, expression, facial shape, and skin tone with less than 0.4% of the original pixel count. This capability is achieved by designing  $\mathbf{O}$  to selectively retrieve relevant features from the target domain, enabling focused extraction of critical characteristics for facial identification while preserving landmarks embedded in the underlying distribution learned by the pre-trained generative model  $\mathcal{G}$ . The model  $\mathcal{G}$  was trained exclusively on face images (FFHQ  $256 \times 256$ ), which is the same training data used for all comparison methods. This targeted feature retrieval, along with compressing the target space into a flat latent structure, enables the reconstruction process not only to interpolate under-sampled data but also to generate outputs closely resembling the original, detailed facial images.

The compression ratios shown in Fig. 4 correspond to measurement levels of 512, 256, 128, 64, 32, and 16, with quantitative assessments provided in Table 1. Crucially, only 64 measurements are already sufficient for accurate recognition with ArcFace [10] and Dlib [30], with accuracy declining significantly below 32 samples. However, even for these extreme compression ratios, critical facial features—such as facial structure, gaze direction, and expression—are remarkably preserved. This underscores the advantage of targeting the latent space of generative models over achieving pixel-perfect accuracy. Beyond visualization, the linear boundaries within the latent space offer significant benefits for downstream applications, enabling effective feature recognition with minimal adjustments.

## 8. Prototype Results

Real world experiments are conducted using our prototype and the pipeline described in Sec. 5. A set of 100 face images is randomly selected from CelebAHQ test set and displayed on a high-dynamic-range Eizo CG3145 monitor, then captured by the single-pixel setup. Fig. 5 illustrates the reconstruction results across three different compression ratios. These same measurements are further utilized by the linear projectors to produce attributes classification, segmentation masks and landmark points for compression ratios of 1:512, 1:2048 and 1:8192 (see Figs. 6, 8 and 7), corresponding to 128 and 32 scalars, respectively, thereby validating the feasibility of the LSI framework.




Figure 4. Illustration of reconstructions achieved through the LSI pipeline with bounded and quantized pixels, using compression ratios from 1:128 to 1:4096 in the simulation setting.

Comp. Ratio	VGGFace $\uparrow$	DLib $\uparrow$	FID $\downarrow$	Acc. $\uparrow$	F1 $\uparrow$	NME $\downarrow$
1:128	91.97%	92.74%	27.38	89.07%	70.00%	1.48
1:256	90.98%	92.68%	26.62	89.15%	70.94%	1.43
1:512	89.61%	91.67%	28.66	89.20%	70.25%	1.48
1:1024	81.12%	87.44%	28.79	88.74%	69.18%	1.52
1:2048	54.72%	77.77%	35.89	88.06%	65.81%	1.67
1:4096	27.22%	59.21%	46.18	86.44%	60.63%	2.01
1:8192	N/A	N/A	N/A	83.81%	53.77%	2.18
1:16384	N/A	N/A	N/A	81.75%	46.36%	2.47

Table 1. Latent Space Imaging is quantitatively evaluated for simulated results across multiple downstream tasks under highly compressed settings. For the image reconstruction task, validation is performed using both the VGGFace [10] and DLib [30] image recognition pipelines. Performance metrics include Fréchet Inception Distance (FID), classification accuracy, F1-mean score for segmentation tasks, and Normalized Mean Error (NME), normalized by the diagonal of the face bounding box. In cases of extreme compression ratios, while recognition metrics and reconstruction may not be applicable (N/A), downstream applications remain feasible.

In the following, we demonstrate how our sensing matrix, optimized for the latent space, effectively supports a wide range of downstream computer vision tasks *without* requiring full image reconstruction. All results presented are based on *real-world experiments* conducted using the prototype described in Sec. 5.

Each application is evaluated on the CelebAHQ [25] dataset, following the architectures and feature projections detailed in Sec. 4. For attribute classification, we use the original 40 labels per image as ground truth. In landmark

detection, we leverage FaRL [72] as a state-of-the-art model to compute 68 landmark points, using the original image  $I$  as the reference. For segmentation tasks, we empirically found that generating masks using reconstructed faces  $I'$  with FaRL as ground truth achieves superior accuracy, given the spatial precision needed for effective segmentation.

### 8.1. Attributes

The classification accuracy of 40 binary attributes is evaluated across the full range of compression ratios. Notably,




Figure 5. Face reconstructions from the experimental setup across different compression ratios.

we achieve 80% precision even at a 1:16384 compression ratio, which corresponds to only 4 scalar measurements. Fig. 6 visualizes the sigmoid-transformed logits after linearly projecting the latent space,  $P_A(\mathbf{L})$ . Despite the decline in similarity observed in Table 1 and Fig. 4, the selected attribute ratios remain close to the original, offering an accurate representation of the original face.

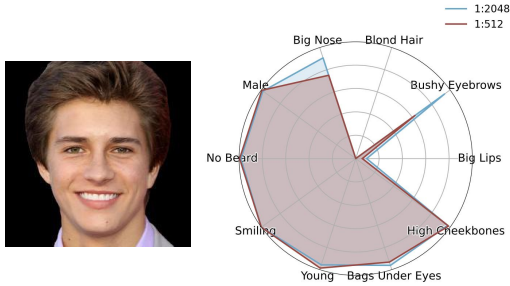


Figure 6. The figure illustrates the sigmoid-transformed logits for compression ratios of 1:256 and 1:2048, overlaid for comparison. These correspond to 256 and 32 scalars, respectively.

## 8.2. Facial Landmarks

Fig. 7 highlights landmark detection acquired by the linear projection  $P_L(\mathbf{L})$  on the same two images across various compression levels, demonstrating high accuracy even in extremely compressed scenarios, such as 1:8192 (8 scalars), compared with the pseudo ground-truth ( $gt$ ) generated by FaRL [72]. The third row shows a non-frontal view of a face, yet the predictions remain close to the reference, indicating robustness.

The quantitative evaluation in Table 1 uses the Normalized Mean Error (NME) metric (Eq. 4), scaled by  $\alpha = 100$

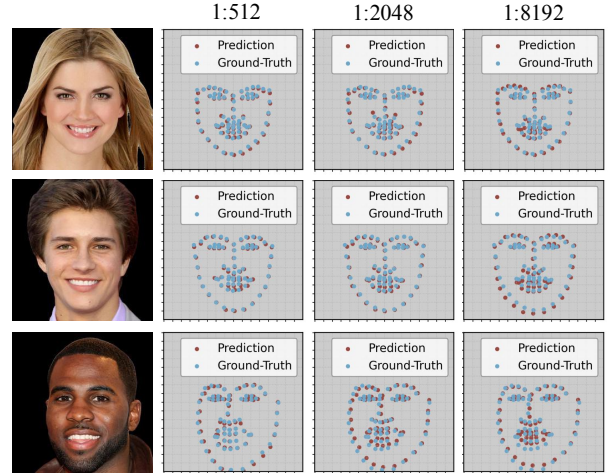


Figure 7. Overlay of predicted and ground-truth positions for 68 facial landmarks across different compression ratios.

for visualization purposes. NME is calculated as the average Euclidean distance between predicted and ground-truth landmarks, normalized by the diagonal of the face bounding box ( $d_{i}^{dg}$ ) to account for variations in face size.

$$\text{NME}_{dg} = \frac{\alpha}{N} \sum_{i=1}^N \frac{\|P_L(\mathbf{L})_i - gt_i\|_2}{d_i^{dg}}. \quad (4)$$

## 8.3. Face Segmentation

The segmentation task is performed using the FaRL model, trained on the CelebAHQ dataset, which includes 18 classes. During our linear projection,  $P_S$ , we aim to learn the entirety of these classes. However, due to the limitations of our generative model, reproducing high-frequency details and unique objects are challenging; for instance, accessories like earrings are not well segmented. Despite this, for general facial features such as nose, hair, face, mouth, ears, and eyebrows, the linear projection of  $\mathcal{G}$  internal features performs well, even under highly compressive settings.

Fig. 8 illustrates segmentation results across a range of different hair styles and different facial hair, showing that the segmentation masks remain consistent despite a four times increase in compression ratio. Additionally, Table 1 highlights the stability of the F1 score across different compression ratios, demonstrating reliable segmentation performance.

Extended results and discussions for all compression ratios, downstream tasks, evaluation metrics, detailed descriptions of the physical setup are provided in the Supplement.






Figure 8. Illustration of face segmentation for two compression levels (1:512 and 1:2048) as well as the pseudo ground-truth for extract from the original image utilizing the FaRL model.

## 9. Comparison with Alternative Approaches

Although our approach diverges from current strategies in compressed imaging, we selected notable models for comparison. Specifically, we included representative models from block-based sensing matrices [15, 54], unfolding networks [60], traditional single-pixel methods with deep learning-based reconstruction [14, 59], Fourier pattern single-pixel methods [71], and optimization-based methods using generative models as priors [22]. Table 2 and Fig. 9 illustrate the comparison quantitative and qualitative simulated results for deep learning single-pixel using Fourier features [71] and unfolding networks [60]. For fair comparison, the sensing matrix for the unfolding network is 1 bit and 8 bit quantized which is required for a plausible physical implementation. Quantization training follows the same strategy adopted by LSI. Detailed methodologies and all mentioned comparisons are elaborated in the Supplement.

Fig. 9 highlights the details of our model which contrasts with traditional compressive sensing strategies, where limited measurements lead to reconstruction algorithms over-smoothing the images, resulting in the loss of critical identity features such as the shape of eyes, facial hair, and other distinguishing landmarks. Table 2 also highlights that even with better quantization deep unfolding methods still outperformed by LSI.

## 10. Discussion and Conclusion

Compared to other imaging techniques, Latent Space Imaging offers substantial benefits in domain-specific scenarios with tight requirements for hardware complexity, memory and bandwidth requirements, or frame rate. As mentioned in the introduction, LSI is strongly inspired by biological




Figure 9. Qualitative comparison between different methods.

	VGGFace $\uparrow$	Dlib $\uparrow$	FID $\downarrow$
FSI-DL	3.30%	13.15%	123.5
SAUNet	17.84%	50.12%	75.53
SAUNet $^\dagger$	39.47%	65.97%	65.79
LSI	<b>90.98%</b>	<b>92.68%</b>	<b>26.62</b>

Table 2. Quantitative comparison between different compressed imaging methods for a 1:256 ratio. SAUNet $^\dagger$  represents a version retrained on 8 bits quantized sensing matrix, while the SAUNet was 1 bit quantized, similar to LSI and FSI-DL.

vision, particularly the human visual system. The relationship between LSI and biological vision is discussed in more detail in the Supplement.

LSI inherits the advantages of a well-trained generative model specialized in a specific domain. Consequently, its encoding learns to represent the most meaningful and essential information within this domain. This capability enables high-quality measurements with high compression ratios. However, it also inherently possesses the limitations of the underlying generative model and the inversion network. Therefore, the model may fail to reconstruct very specific features not present during training, such as hats, glasses, makeup, or other distinct image details, and may struggle to generalize to a multi-domain approach.

Nonetheless, we believe that LSI is a promising new paradigm for domain-specific imaging systems, including privacy-preserving cameras that nonetheless excel at a specific imaging task. LSI is also a promising avenue for imaging systems ultra-low hardware requirements (both in terms of optical/electronic complexity and compute power), and/or ultra-high frame rate imaging. These applications will require alternative hardware implementations compared to our single-pixel camera prototype, which was chosen primarily due to easy reconfigurability. Several concepts for such alternative hardware designs are discussed in the Supplement.



## References

- [1] Rameen Abdal, Yipeng Qin, and Peter Wonka. Image2StyleGAN: How to embed images into the StyleGAN latent space? In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2019. 3
- [2] Yochai Blau and Tomer Michaeli. The perception-distortion tradeoff. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6228–6237, 2018. 5
- [3] Ashish Bora, Ajil Jalal, Eric Price, and Alexandros G Dimakis. Compressed sensing using generative models. In *International conference on machine learning*, pages 537–546. PMLR, 2017. 2
- [4] Ayan Chakrabarti. Learning sensor multiplexing design through back-propagation. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 3089–3097, Red Hook, NY, USA, 2016. Curran Associates Inc. 2
- [5] Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10193–10202, 2019. 2
- [6] Julie Chang, Vincent Sitzmann, Xiong Dun, Wolfgang Heidrich, and Gordon Wetzstein. Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification. *Sci. Rep.*, 8(1):12324, 2018. 2
- [7] Yunje Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020. 5
- [8] Ilya Chugunov, Seung-Hwan Baek, Qiang Fu, Wolfgang Heidrich, and Felix Heide. Mask-tof: Learning microlens masks for flying pixel correction in time-of-flight imaging. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9116–9126, 2021. 2
- [9] Edo Collins, Raja Bala, Bob Price, and Sabine Susstrunk. Editing in style: Uncovering the local semantics of gans. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5771–5780, 2020. 3
- [10] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4685–4694, 2019. 5, 6
- [11] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4690–4699, 2019. 5
- [12] D.L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006. 2
- [13] Marco F Duarte, Mark A Davenport, Dharmpal Takhar, Jason N Laska, Ting Sun, Kevin F Kelly, and Richard G Baraniuk. Single-pixel imaging via compressive sampling. *IEEE Signal Processing Magazine*, 25(2):83–91, 2008. 2
- [14] Fabio Ferri, D Magatti, LA Lugiato, and A Gatti. Differential ghost imaging. *Physical review letters*, 104(25):253603, 2010. 8
- [15] Hongping Gan, Minghe Shen, Yi Hua, Chunyan Ma, and Tao Zhang. From patch to pixel: A transformer-based hierarchical framework for compressive image sensing. *IEEE Transactions on Computational Imaging*, 9:133–146, 2023. 8
- [16] Harel Haim, Shay Elmaleh, Raja Giryes, Alex M Bronstein, and Emanuel Marom. Depth estimation from a single image using deep learned phase coded mask. *IEEE Transactions on Computational Imaging*, 4(3):298–310, 2018. 2
- [17] Lei He, Guanghui Wang, and Zhanyi Hu. Learning depth from single images with deep neural network embedding focal length. *IEEE Transactions on Image Processing*, 27(9):4676–4689, 2018. 2
- [18] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Adv. Neural Inform. Process. Syst.*, 30, 2017. 5
- [19] Catherine F Higham, Roderick Murray-Smith, Miles J Padgett, and Matthew P Edgar. Deep learning for real-time single-pixel video. *Sci. Rep.*, 8(1):2369, 2018. 2
- [20] Wenxin Huang, Fei Wang, Xiangyu Zhang, Ying Jin, and Guohai Situ. Learning-based adaptive under-sampling for fourier single-pixel imaging. *Opt. Lett.*, 48(11):2985–2988, 2023. 2
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks. *Adv. Neural Inform. Process. Syst.*, 29, 2016. 5
- [22] Ajil Jalal, Liu Liu, Alexandros G Dimakis, and Constantine Caramanis. Robust compressed sensing using generative models. *Advances in Neural Information Processing Systems*, 33:713–727, 2020. 2, 8
- [23] Jost B Jonas, Andreas M Schmidt, JA Müller-Bergh, UM Schlötzer-Schrehardt, and GO Naumann. Human optic nerve fiber count and optic disc size. *Investigative Ophthalmology & Visual Science*, 33(6):2012–2018, 1992. 1
- [24] Maya Kabkab, Pouya Samangouei, and Rama Chellappa. Task-aware compressed sensing with generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018. 2
- [25] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Int. Conf. Learn. Represent.*, 2018. 5, 6
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks, 2019. 2, 5
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of StyleGAN. In *Proc. CVPR*, 2020.
- [28] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34:852–863, 2021. 2
- [29] Ori Katz, Yaron Bromberg, and Yaron Silberberg. Compressive ghost imaging. *Applied Physics Letters*, 95(13), 2009. 2
- [30] Vahid Kazemi and Josephine Sullivan. One millisecond face alignment with an ensemble of regression trees. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1867–1874, 2014. 5, 6

- [31] Varun A Kelkar and Mark A Anastasio. Prior image-constrained reconstruction using style-based generative models. *arXiv preprint arXiv:2102.12525*, 2021. 2
- [32] Varun A Kelkar, Sayantan Bhadra, and Mark A Anastasio. Compressible latent-space invertible networks for generative model-constrained image reconstruction. *IEEE transactions on computational imaging*, 7:209–223, 2021. 2
- [33] Michael Kellman, Emrah Bostan, Michael Chen, and Laura Waller. Data-driven design for Fourier ptychographic microscopy. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–8, 2019. 2
- [34] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4015–4026, 2023. 5
- [35] Jeremy Klotz and Shree K Nayar. Minimalist vision with freeform pixels. In *European Conference on Computer Vision*, pages 329–346. Springer, 2025. 2
- [36] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 4681–4690, 2017. 5
- [37] Xing Lin, Yair Rivenson, Nezh T Yardimci, Muhammed Veli, Yi Luo, Mona Jarrahi, and Aydogan Ozcan. All-optical machine learning using diffractive deep neural networks. *Science*, 361(6406):1004–1008, 2018. 2
- [38] Hanxiao Liu, Zihang Dai, David So, and Quoc V Le. Pay attention to MLPs. *Adv. Neural Inform. Process. Syst.*, 34: 9204–9215, 2021. 3
- [39] Fabian Mentzer, David Minnen, Eirikur Agustsson, and Michael Tschannen. Finite scalar quantization: VQ-VAE made simple. In *The Twelfth International Conference on Learning Representations*, 2024. 5
- [40] Robert S Molday and Orson L Moritz. Photoreceptors at a glance. *Journal of Cell Science*, 128(22):4039–4045, 2015. 1
- [41] Chong Mou, Qian Wang, and Jian Zhang. Deep generalized unfolding networks for image restoration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 17399–17410, 2022. 2
- [42] Elias Nehme, Daniel Freedman, Racheli Gordon, Boris Ferdman, Lucien Weiss, Onit Alalouf, Tal Naor, Reut Orange, Tomer Michaeli, and Yoav Shechtman. DeepSTORM3D: dense 3D localization microscopy and PSF design by deep learning. *Nat. Methods*, 17:1–7, 2020. 2
- [43] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015. 5
- [44] Parita Pooj, Michael Grossberg, Peter N Belhumeur, and Shree K Nayar. The minimalist camera. In *BMVC*, page 141, 2018. 2
- [45] Gang Qu, Ping Wang, and Xin Yuan. Dual-scale transformer for large-scale single-pixel imaging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25327–25337, 2024. 2
- [46] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with vq-vae-2, 2019. 5
- [47] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a StyleGAN encoder for image-to-image translation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2021. 3
- [48] Saad Rizvi, Jie Cao, Kaiyu Zhang, and Qun Hao. Improving imaging quality of real-time fourier single-pixel imaging via deep learning. *Sensors*, 19(19), 2019. 2
- [49] Axel Sauer, Katja Schwarz, and Andreas Geiger. StyleGAN-xl: Scaling StyleGAN to large diverse datasets. In *ACM SIGGRAPH 2022 conference proceedings*, pages 1–10, 2022. 2, 3
- [50] Sefik Ilkin Serengil and Alper Ozpinar. Lightface: A hybrid deep face recognition framework. In *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)*, pages 1–5, 2020. 5
- [51] Yoav Shechtman, Lucien Weiss, Adam Backer, Maurice Lee, and William Moerner. Multicolour localization microscopy by point-spread-function engineering. *Nat. Photonics*, 10, 2016. 2
- [52] Yujun Shen, Jinjin Gu, Xiaou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9243–9252, 2020. 3
- [53] Vincent Sitzmann, Steven Diamond, Yifan Peng, Xiong Dun, Stephen Boyd, Wolfgang Heidrich, Felix Heide, and Gordon Wetzstein. End-to-end optimization of optics and image processing for achromatic extended depth of field and super-resolution imaging. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 2
- [54] Jiechong Song, Chong Mou, Shiqi Wang, Siwei Ma, and Jian Zhang. Optimization-inspired cross-attention transformer for compressive sensing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 6174–6184, 2023. 2, 8
- [55] Kai Song, Yaoxing Bian, Ku Wu, Hongrui Liu, Shuangping Han, Jiaming Li, Jiazhao Tian, Chengbin Qin, Jianyong Hu, and Liantuan Xiao. Single-pixel imaging based on deep learning, 2023. 2
- [56] Qilin Sun, Ethan Tseng, Qiang Fu, Wolfgang Heidrich, and Felix Heide. Learning Rank-1 diffractive optics for single-shot high dynamic range imaging. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1386–1396, 2020. 2
- [57] Qilin Sun, Congli Wang, Qiang Fu, Xiong Dun, and Wolfgang Heidrich. End-to-end complex lens design with differentiate ray tracing. *ACM Trans. Graph.*, 40(4), 2021. 2
- [58] Aaron van den Oord, Oriol Vinyals, and Koray Kavukcuoglu. Neural discrete representation learning, 2018. 5
- [59] Fei Wang, Chenglong Wang, Chenjin Deng, Shensheng Han, and Guohai Situ. Single-pixel imaging using physics enhanced deep learning. *Photon. Res.*, 10(1):104–110, 2022. 8
- [60] Ping Wang and Xin Yuan. Saunet: Spatial-attention unfolding network for image compressive sensing. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 5099–5108, 2023. 2, 8

- [61] Ping Wang, Yulun Zhang, Lishun Wang, and Xin Yuan. Hierarchical separable video transformer for snapshot compressive imaging. *arXiv preprint arXiv:2407.11946*, 2024.
- [62] Xiaoyang Wang and Hongping Gan. Ufc-net: Unrolling fixed-point continuous network for deep compressive sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25149–25159, 2024. [2](#)
- [63] Kaixuan Wei, Xiao Li, Johannes Froech, Praneeth Chakravarthula, James Whitehead, Ethan Tseng, Arka Majumdar, and Felix Heide. Spatially varying nanophotonic neural networks. *Sci. Adv.*, 10(45), 2024. [2](#)
- [64] Gordon Wetzstein, Aydogan Ozcan, Sylvain Gigan, Shan-hui Fan, Dirk Englund, Marin Soljačić, Cornelia Denz, David AB Miller, and Demetri Psaltis. Inference in artificial intelligence with deep optics and photonics. *Nature*, 588(7836):39–47, 2020. [2](#)
- [65] Yicheng Wu, Vivek Boominathan, Huaijin Chen, Aswin Sankaranarayanan, and Ashok Veeraraghavan. Phase-CAM3D—learning phase masks for passive single view depth estimation. In *2019 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2019. [2](#)
- [66] Weihao Xia, Yulun Zhang, Yujiu Yang, Jing-Hao Xue, Bolei Zhou, and Ming-Hsuan Yang. GAN inversion: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2022. [3](#)
- [67] Jianjin Xu and Changxi Zheng. Linear semantics in generative adversarial networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9351–9360, 2021. [3](#)
- [68] Xu Yang, Pengfei Jiang, Mingfeng Jiang, Lu Xu, Long Wu, Chenghua Yang, Wei Zhang, Jianlong Zhang, and Yong Zhang. High imaging quality of fourier single pixel imaging based on generative adversarial networks at low sampling rate. *Optics and Lasers in Engineering*, 140:106533, 2021. [2](#)
- [69] Jian Zhang and Bernard Ghanem. Ista-net: Interpretable optimization-inspired deep network for image compressive sensing. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 1828–1837, 2018. [2](#)
- [70] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 586–595, 2018. [5](#)
- [71] Zibang Zhang, Xiao Ma, and Jingang Zhong. Single-pixel imaging by means of fourier spectrum acquisition. *Nat. Commun.*, 6(1):6225, 2015. [2](#), [8](#)
- [72] Yinglin Zheng, Hao Yang, Ting Zhang, Jianmin Bao, Dongdong Chen, Yangyu Huang, Lu Yuan, Dong Chen, Ming Zeng, and Fang Wen. General facial representation learning in a visual-linguistic manner. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 18697–18709, 2022. [6](#), [7](#)
- [73] Yang Zhou, Zichong Chen, and Hui Huang. Deformable one-shot face stylization via dino semantic guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7787–7796, 2024. [5](#)