

UPLiFT: Efficient Pixel-Dense Feature Upsampling with Local Attenders

Matthew Walmer¹ Saksham Suri² Anirud Aggarwal¹ Abhinav Shrivastava¹
¹University of Maryland, College Park ²Meta

Abstract

The space of task-agnostic feature upsampling has emerged as a promising area of research to efficiently create denser features from pre-trained visual backbones. These methods act as a shortcut to achieve dense features for a fraction of the cost by learning to map low-resolution features to high-resolution versions. While early works in this space used iterative upsampling approaches, more recent works have switched to cross-attention-based methods, which risk falling into the same efficiency scaling problems of the backbones they are upsampling. In this work, we demonstrate that iterative upsampling methods can still compete with cross-attention-based methods; moreover, they can achieve state-of-the-art performance with lower inference costs. We propose **UPLiFT**, an architecture for **Universal Pixel-dense Lightweight Feature Transforms**. We also propose an efficient **Local Attender** operator to overcome the limitations of prior iterative feature upsampling methods. This operator uses an alternative attentional pooling formulation defined fully locally. We show that our Local Attender allows UPLiFT to maintain stable features throughout upsampling, enabling state-of-the-art performance with lower inference costs than existing pixel-dense feature upsamplers. In addition, we apply UPLiFT to generative downstream tasks and show that it achieves competitive performance with state-of-the-art Coupled Flow Matching models for VAE feature upsampling. Altogether, UPLiFT offers a versatile and efficient approach to creating denser features.

1. Introduction

With the increasing spread of deep learning networks in various real-world applications, efficiency for visual information processing continues to be of key importance. General-purpose, pre-trained visual backbones, such as DINO [4, 34, 42], can be a driving force for efficiency, as these models provide a strong starting point for various visual applications, and reduce the total training cost needed to develop systems for new tasks. However, DINO and other Vision Transformer (ViT) [13] backbones face a fundamental problem: they must down-sample the spatial

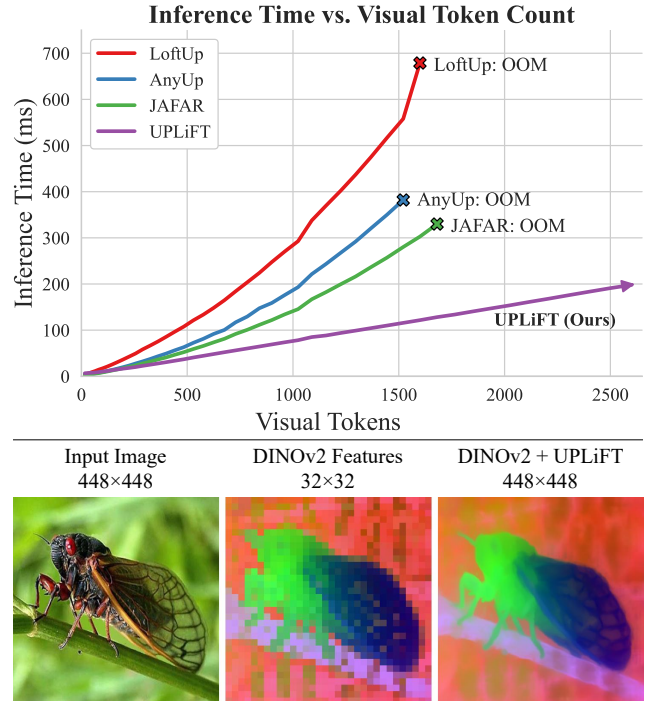


Figure 1. **UPLiFT time-scaling and dense features.** We present **UPLiFT**, an efficient feature-upsampler that leverages our new **Local Attender** to extract semantically-stable, pixel-dense features. (Top) UPLiFT’s inference time and memory scales linearly with the number of visual tokens, while other recent SOTA methods face quadratic scaling. (Bottom) PCA visualization of low-resolution DINOv2 features and pixel-dense UPLiFT features.

dimension to make visual tokens, which limits the final feature map density. For many applications, denser features are a necessity for effective performance, and while it is possible to increase the token density of ViTs to yield denser features [2], the cost of self-attention operations scale quadratically with the number of visual tokens in the image. Thus, we face a fundamental trade-off between the benefits of denser features and the computational cost they incur.

For this reason, feature super-resolution methods have recently grown in popularity [10, 15, 18, 43, 51]. These methods act as task-agnostic add-ons to existing visual feature extracting backbones to directly transform their coarse, low-resolution features into high-resolution ones that pre-

serve the original semantics. Such approaches have the potential to combine the benefits of dense features with the strength of pre-trained visual backbones, while avoiding the quadratic scaling cost of self-attention. Early approaches [15, 43] used multiple steps of iterative $2\times$ feature upsampling to push the coarse backbone features to pixel-scale-density. More recent approaches have switched to using cross-attention-based feature pooling to directly upsample to any desired output size. While these approaches offer strength and flexibility, they risk falling back into the pitfall of quadratic time scaling.

In this work, we examine the recent advances in task-agnostic feature super-resolution methods, and propose a new, efficient upsampler architecture called **UPLiFT** (Universal Pixel-dense Lightweight Feature Transforms). We draw inspiration from LiFT [43], and we demonstrate that iterative upsampling methods are still a strong competitor with more recent cross-attention-based approaches. Not only does our UPLiFT model achieve state-of-the-art performance on several dense prediction tasks, it does so with lower inference costs than comparable recent methods. To achieve this, we propose a new visual operator based on the observations of [48], which showed that some ViT attention heads consistently learn to compute offset local attentional operations. As such, we propose a new **Local Attender** operator, an alternative formulation for attention which eschews the now standardized Query-Key-Value formulation and instead defines all operations based on local relative positions. We show that this localized attentional pooling mechanism achieves the same benefits of consistent feature semantics while avoiding the scaling costs of cross-attention. As shown in Figure 1, our UPLiFT method maintains roughly linear time scaling with respect to visual token count, while recent state-of-the-art cross-attention-based methods display quadratic time and high memory use.

Finally, we apply our UPLiFT model to a range of dense downstream tasks, including both predictive and generative tasks. In addition to pre-trained ViT features, we apply UPLiFT to upsample the latent features of Variational Autoencoders (VAE) [21], enabling applications in image super-resolution and efficient image generation. We show that in this domain, UPLiFT achieves competitive performance with state-of-the-art Coupled Flow-Matching (CFM) models [41], while running with a much lower inference cost. Overall, our contributions include:

- **UPLiFT**, an efficient latent feature upsampler that can be applied to create high-quality, pixel-dense features.
- A **Local Attender** operator that reformulates attentional pooling into an efficient, locally-defined operation. We show that our Local Attender enables UPLiFT to maintain consistent feature semantics, while avoiding the higher costs of comparable recent approaches.
- State-of-the-art performance on several dense predictive

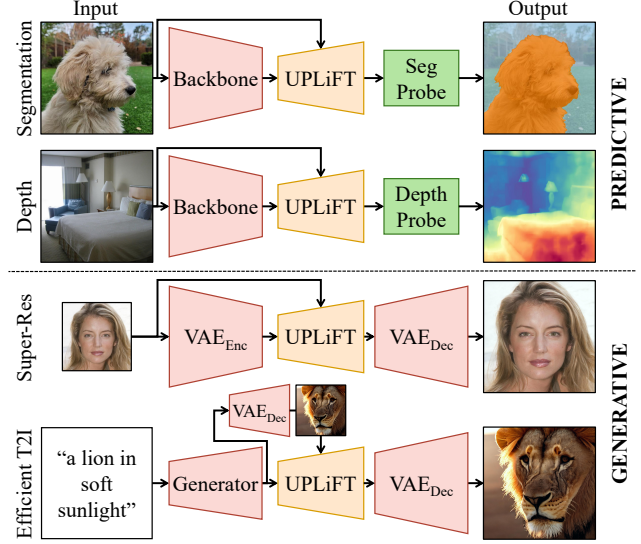


Figure 2. **UPLiFT Tasks.** We demonstrate our UPLiFT feature upsampler for applications in both predictive and generative tasks. This includes semantic segmentation, monocular depth estimation, image super-resolution, and efficient text-to-image generation.

tasks, surpassing prior feature upsamplers while having faster inference speeds.

- An extended study of feature upsampling in the realm of generative tasks, showing the strength and efficiency of UPLiFT in comparison to a more computationally expensive Coupled Flow Matching method.

2. Related Works

2.1. Feature Upsampling for Visual Backbones

Higher-resolution visual features can be beneficial for many tasks, but extracting denser features comes with a larger computational cost. Feature-upsampling methods are a shortcut to create denser features from pre-trained, general-purpose visual backbones like DINO [4, 34, 42], MoCo [5, 6, 16], MAE [17], CLIP [37], and SigLip [45, 53]. In the past, many works have been proposed to design learnable feature adapters which are trained with supervision from a particular downstream task [29, 31, 49, 55]. More recently, there has been an increasing number of methods proposed to output denser features in a task-agnostic way [10, 15, 18, 43, 51]. These methods are typically self-supervised, and are usually trained to upsample the features of a particular pre-trained backbone. Such upsamplers can streamline development for dense tasks, by providing semantically rich and dense features directly out of the box.

Early works in this area [15, 43] used iterative upsampling with simple modules to create denser features. FeatUp [15] proposed two feature upsampling methods, the first using an improved Joint Bilinear Upsampler (JBU) [22] and the second using per-image Implicit Networks [8]. The first

method provides high-speed upsampling but lower upsampling quality, while the latter provides incredibly sharp and high-quality upsampled features, but has a high inference cost as it requires training an implicit model for each image. LiFT [43] proposed a Lightweight Feature Transforming network with a simple self-supervised learning process to create $2\times$ upsampled features, with the potential for larger upsampling through iterative application of the same module. While this approach is simple and efficient, it has been shown that LiFT’s iterative upsampling can lead to semantic drift and degraded features [10]. In this work, we propose a new iterative feature upsampler, UPLiFT, which overcomes the limitations of prior similar approaches.

More recent works [10, 18, 51] have shifted to using cross-attention-based feature resampling. LoftUp [18] and JAFAR [10] both propose a query-key-value (QKV) cross-attention [47] approach that combines high-resolution queries derived from an input image, with low-resolution keys and values derived from backbone features to perform flexible and direct upsampling. Moreover, this approach acts as an effective regularizer, ensuring the output features maintain a similar distribution to the input features. AnyUp [51] further adapts the JAFAR architecture [10] into a backbone-agnostic model, with a special adapter layer that allows it to generalize to other feature extractors at inference time. These approaches show strong results, but their use of QKV cross-attention makes them less efficient, giving them quadratic time-scaling with respect to the number of visual tokens. In contrast, our proposed UPLiFT architecture enjoys the same benefits of stable features, while also remaining linear in time with respect to the number of image patches. We also recognize that UPLiFT’s use of fixed-step upsamplers makes it slightly less flexible than cross-attention-based methods; however, given that most approaches simply aim to upsample to pixel-dense features, we believe this is not a significant detriment.

2.2. Feature Upsampling for Generative Tasks

The majority of prior works in task-agnostic feature upsampling have focused on visual feature extractors typically designed for predictive downstream tasks. Meanwhile, relatively little attention has been given to designing upsamplers for features for generative tasks. For example, the feature-spaces of Variational Autoencoders [21] are now commonly used in Latent Diffusion models [39] to provide a compressed space for image generation. Because VAEs have decoders, upsampled VAE features can easily be decoded to enable tasks like image super-resolution and efficient text-to-image generation. We note that both of these tasks are areas of significant research in their own right. For example, many methods exist that train primarily end-to-end generative architectures for super-resolution [7, 11, 24, 33, 40]. However, for this work, we specifically

wish to study methods which upsample VAE latent features, with image-super-resolution being a possible application of said features. Meanwhile, extensive research has also been devoted to adapting pre-trained visual diffusion models to generate images at resolutions larger than their original training scale [19, 25, 27, 36, 44, 52]. These methods may achieve a similar effect: adapting a pre-trained generator to create larger images, but they involve modifying the generator module/process itself, not the generated latents. These approaches also typically come with significantly increased generation costs, which feature upsampling could avoid. Moreover, recent generator architectures like [23] are designed to flexibly output features at a range of latent code sizes, but they still face the trade-off that visual feature extractors face: larger features come with a larger extraction/generation cost.

For this work, we aim to compare with methods that specifically learn to upsample VAE latent features. The state-of-the-art method in said space is Coupled Flow-Matching (CFM) [41], which uses a powerful Flow-Matching model [1, 28, 30] conditioned on low-resolution VAE features to predict high-resolution VAE features. CFM is trained using a learning objective very similar to LiFT and UPLiFT, and it can be applied to upsample VAE latent codes either from encoded real images or from a generator like Stable Diffusion [38]. Unlike UPLiFT, CFM performs upscaling by first decoding the latent representation into pixel space, applying bilinear upscaling therein, and then re-encoding the result into latent space. The CFM module then refines the high-resolution features; however, the overall effect is still to produce high resolution latents from low resolution ones. In this work, we show that UPLiFT can also learn to effectively upsample VAE features that produce output images with similar visual fidelity to CFM. Moreover, UPLiFT achieves this with a fraction of the parameters, training data, and inference cost of CFM.

3. UPLiFT Approach

In this work, we propose an iteratively-growing feature upsampling approach inspired by LiFT [43]. While LiFT achieved effective $2\times$ feature upsampling, its approach to iterative upsampling for pixel-dense features could lead to semantic drift and degraded performance in downstream tasks. In this work, we present UPLiFT, a new upsampler architecture that overcomes these issues to efficiently produce pixel-dense features while maintaining consistency in the feature distribution. Thanks to our Local Attender, UPLiFT’s speed and memory requirements scale linearly with respect to the number of input visual patches/tokens, which leads to better speed and scaling properties than recent cross-attention-based methods.

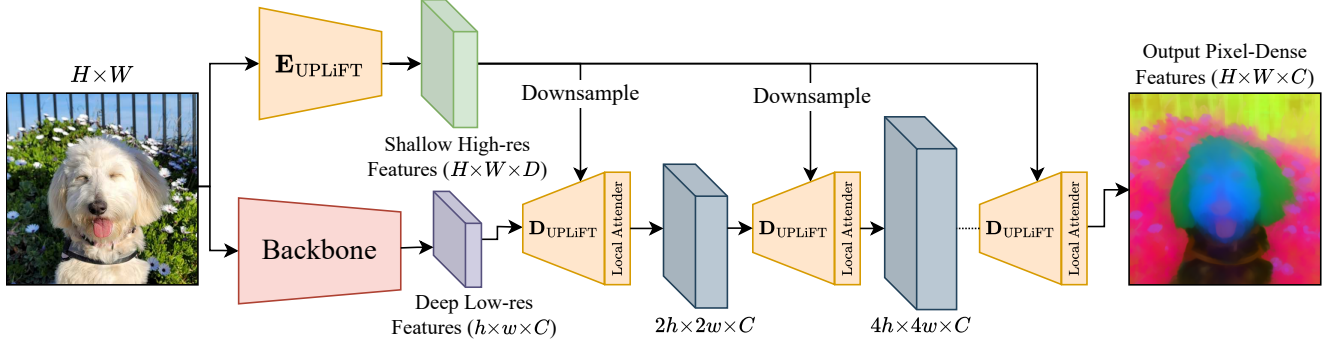


Figure 3. **UPLiFT Inference.** At inference time, our UPLiFT Encoder (E_{UPLiFT}) produces shallow but dense features to guide all subsequent upsampling steps. Iterative application of the UPLiFT Decoder (D_{UPLiFT}) upsamples the low-resolution backbone features to pixel-density. Our proposed Local Attender module is integrated with the UPLiFT Decoder to maintain iterative feature consistency.

3.1. Architecture Overview

We first illustrate our UPLiFT architecture in its inference-time configuration, as shown in Figure 3. UPLiFT has two primary modules: the UPLiFT encoder, denoted as E_{UPLiFT} , and the UPLiFT decoder, denoted as D_{UPLiFT} . D_{UPLiFT} does the primary work of predicting upsampled features. Both are simple convolutional modules, which use strided convolution for downsampling or transpose convolution for upsampling. Note that we train a single compact D_{UPLiFT} that performs $2\times$ upsampling, and that same module is applied multiple times to achieve pixel-dense upsampling. Meanwhile, this decoder is also guided by additional high-resolution features extracted by E_{UPLiFT} from the input image. We follow the same intuition as LiFT: the input image already is a strong source of high-resolution information to guide feature upsampling. However, unlike LiFT, which required re-running its image encoder on every step with bilinearly upsampled input images, we instead design E_{UPLiFT} such that it only runs once with the base input image. To do this, we make E_{UPLiFT} into a shallow but dense encoder which preserves the original pixel density in its output features. These dense encoder features are downsampled with nearest-neighbors downsampling to the appropriate size to guide each upsampling step with D_{UPLiFT} . These design choices make UPLiFT both efficient and compact. However, as shown by [10], this process of iterative upsampling risks introducing semantic drift, which we address in the following section with our proposed Local Attender.

3.2. Local Attender

Preserving semantic consistency in the upsampled features is key to maintaining the strength of their representations. Recent works like [18] and [10] use cross attention from high-resolution queries to low-resolution keys and values to perform upsampling. This approach implicitly acts as a strong form of feature regularization, as it enforces a constraint that the upsampled features must be a linear combination of the input features. However, these cross-attention-

based approaches lead to quadratic-scaling costs. Moreover, we hypothesize that full cross attention over all features is not necessary, and information about only local features should be sufficient to upsample the features in a given image region. We note that the recent work AnyUp [51] shares this hypothesis, and modifies JAFAR by introducing windowed attention. However, as shown by Figure 1, AnyUp’s windowed attention still suffers from quadratic time scaling.

Instead, we propose a new formulation for a local attention operator, which achieves the same feature-regularization benefits of cross-attention while maintaining efficiency and linear-time-scaling. We draw inspiration from [48], which showed that self-attention heads in ViTs often learn to attend to local positions with a fixed directional offset. Based on this, we propose **Local Attender**, which instead operates over a pre-set local neighborhood defined by a collection of fixed directional offsets. Under our Local Attender design, all attentional operations are defined relative to the current token’s position, removing the need for positional embeddings. We illustrate our Local Attender operator in Figure 4, and define it as follows.

First, we assume two input feature maps, a “Guide Feature” G and a “Value Feature” V . The Local Attender operator uses G to guide a linear resampling of V to produce a new feature map output. Specifically, G is used to predict an attention map over a local neighborhood around each token in V . For an initial case, assume G and V have the same spatial dimensions, $H \times W$, but different channel depths, C_G and C_V . Next, we define a fixed neighborhood \mathbf{N} , which is a set of $2D$ directional integer offsets, (i, j) . For a given visual token $V_{x,y}$ in V , it can only attend to positions $V_{x+i,y+j}$ for $(i, j) \in \mathbf{N}$. In this way, the neighborhood \mathbf{N} defines the range and size of the local attentional operations computed by the Local Attender. Moreover, the size and pattern of \mathbf{N} can be flexibly defined. We present experiments with several neighborhood designs in the Supplemental Material. In general, let $|\mathbf{N}| = n$.

Next, we apply a 1×1 convolutional layer to G to con-

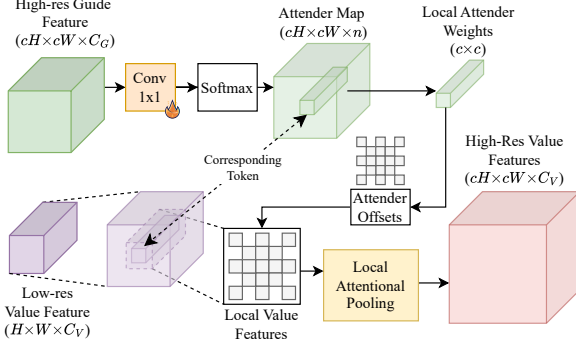


Figure 4. **Local Attender Operator.** We propose a streamlined and efficient local attention operator, which gathers features over a set neighborhood defined by fixed direction offsets.

vert it to shape $H \times W \times n$. This convolutional layer is the only learnable element of the Local Attender. We then apply position-wise softmax to this feature, creating the “Attender Map” A . For a given (x, y) position in A , the value $A_{x,y,k}$ represents the attentional weight to apply to the feature at $V_{x+i_k, y+j_k}$ where (i_k, j_k) is the k^{th} offset in N . In practical terms, we compute a set of offset feature maps by applying the offsets of N to Value Feature V with replication padding. We then multiply these maps by the Attender Map A and we sum along the newly created “neighborhood-dimension” to produce the final locally-attended features of shape $H \times W \times C_V$. Let T equal the number of spatial token in G . Then the memory and compute cost of the Local Attender scales as $O(nT)$, or linearly in T as n is a constant.

Finally, we relax an initial assumption to allow Local Attender to act as an upsampling operator. Instead, assume V has shape $H \times W \times C_V$ and guide G has shape $cH \times cW \times C_G$, where $c \in \mathbb{Z}$. We first group the tokens of G into an $H \times W$ grid of cells, each cell being of shape $c \times c$. For all tokens in a given cell $C_{x,y}$ with $(x, y) \in H \times W$, the neighborhood they can attend to is defined around the corresponding Value token at $V_{x,y}$. As such, the size of the neighborhoods remains the same, but each token in V now corresponds to a $c \times c$ cell in G . Once again, the output features are linear sums of features in V , ensuring feature consistency with respect to V . The output map size is determined by G , with the final output being $cH \times cW \times C_V$.

We integrate our Local Attender operator as the final step of our UPLiFT decoder, with the initial decoder output as G and the original backbone feature as V . We show that the Local Attender ensures the preservation of the backbone feature distribution, which allows UPLiFT to efficiently create semantically-strong, pixel-dense features. In the following section, we show that this leads to state-of-the-art performance with faster inference speeds than recent methods.

3.3. UPLiFT Training

We train UPLiFT with a low-to-high-res self-supervised feature prediction objective based on LiFT, though we adapt

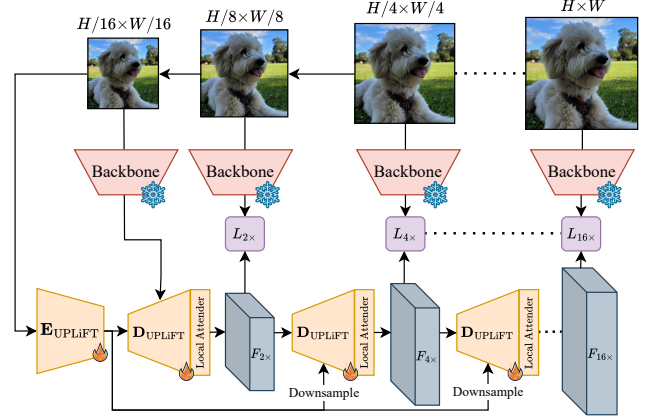


Figure 5. **UPLiFT Training.** UPLiFT uses a multi-step training strategy where the feature reconstruction loss is applied at all intermediate steps. All lift decoders (D_{UPLiFT}) share the same weights.

our training to explicitly include multiple upsampling steps during training. We illustrate our UPLiFT training process in Figure 5. Let B represent the frozen, pre-trained visual backbone. Let E_U and D_U denote the UPLiFT Encoder and Decoder, which are trained jointly in this process. Let $H \times W$ denote the maximum size an image is loaded at. We refer to this high-resolution image as I . We also specify a “training depth” d , which describes the scale of downsampling applied to the images that are input to UPLiFT. For a given depth d , the ground-truth image will be downsampled by a factor of 2^d to size $(H/2^d) \times (W/2^d)$. We denote this low-resolution image as I' . Next, we pass both I and I' through B to extract low-resolution and high-resolution features $F = B(I)$ and $F' = B(I')$. Additionally, I' is passed through E_U to produce features F'_E which are used to guide upsampling. Note that with our dense UPLiFT encoder, F'_E has the same spatial resolution as I' , while F' has lower-resolution with size $(H/2^d p) \times (W/2^d p)$, where p is the patch size.

Next, we perform feature upsampling. To start, D_U receives the low resolution feature F' and guiding feature F'_E which is down-sampled to $2 \times$ the scale of F' . D_U predicts upsampled features with resolution $(H/2^{d-1} p) \times (W/2^{d-1} p)$, which are then provided to the Local Attender to upsample F' . The result is $F'_{2 \times}$, which is two times larger than the original low-resolution features. This process is repeated d times, until we produce $F'_{2^d \times}$ with size $(H/p) \times (W/p)$, the same as F . We then compute the reconstruction loss as the L_2 distance between the upsampled feature and the high-resolution feature:

$$\mathcal{L}_{\text{simple}} = \mathcal{D}_{L_2}(F'_{2^d \times}, F) \quad (1)$$

This simple learning objective explicitly incorporates multiple applications of D_U during training, which helps UPLiFT learn to create pixel-dense features at inference time. However, this training does not necessarily need to be ap-

Table 1. **Segmentation and Depth Estimation.** UPLiFT surpassed all prior upsampling methods in semantic segmentation on four datasets. For depth estimation, it is second best for δ_1 and tied for best in RMSE. UPLiFT achieves these scores with faster inference speeds than recent SOTA methods. Best results are marked in **bold** and second best are underlined.

			Semantic Segmentation								Depth Estimation	
Upsampler			COCO		VOC		ADE20K		Cityscapes		COCO	
Method	Params (M)	Time (ms)	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow	δ_1 \uparrow	RMSE \downarrow
Nearest	–	0.6	56.41	77.11	78.29	94.37	37.86	72.17	54.87	90.56	56.66	0.73
Bilinear	–	2.8	59.41	79.28	81.62	95.44	40.43	74.12	59.71	92.56	58.83	0.68
LiFT-2 \times	1.2	3.8	58.28	78.97	82.46	95.73	39.74	73.79	61.07	93.09	57.17	0.70
LiFT	1.2	51.9	57.42	78.46	80.97	95.37	38.95	73.34	61.98	<u>93.61</u>	55.07	0.73
FeatUp	0.2	109.6	61.77	80.99	83.52	96.06	42.07	75.52	60.50	93.12	60.01	0.66
LoftUp	4.3	223.5	<u>62.19</u>	<u>81.35</u>	<u>84.63</u>	<u>96.33</u>	42.16	75.79	<u>62.09</u>	93.55	58.69	0.68
JAFAR	0.7	111.7	61.71	81.01	84.38	96.22	41.96	75.43	61.89	93.52	60.59	<u>0.65</u>
AnyUp	0.9	146.7	62.08	81.31	84.33	96.23	<u>42.25</u>	<u>75.80</u>	61.33	93.44	61.32	0.63
UPLiFT	0.8	79.4	62.55	81.57	85.21	96.51	42.97	76.00	65.38	94.41	<u>61.16</u>	0.63

plied with $d = \log_2(p)$, as shallower depths can be sufficient to train UPLiFT. In addition, for each D_U upsampling step, an additional intermediate feature map can be extracted to derive an additional loss term. Denote the full set of upsampled feature maps as:

$$\begin{aligned} \mathbf{F}' &= \{F'_{2\times}, F'_{4\times}, \dots, F'_{2^{d-1}\times}, F'_{2^d\times}\} \\ &= \{F'_{1/2^{d-1}}, F'_{1/2^{d-2}}, \dots, F'_{1/2}, F'_{1/1}\} \end{aligned} \quad (2)$$

We also compute intermediate ground truth feature maps by downsampling I to intermediate resolutions, with $F_{1/2^k} = B(I_{1/2^k})$, to give the set of target feature maps:

$$\mathbf{F} = \{F_{1/2^{d-1}}, F_{1/2^{d-2}}, \dots, F_{1/2}, F_{1/1}\} \quad (3)$$

Then we can compute the total loss for depth d with all intermediate feature maps as:

$$\mathcal{L}_d = \sum_{k=1}^d \mathcal{D}_{L2}(F'_{1/2^{d-k}}, F_{1/2^{d-k}}) \quad (4)$$

In practice, we find UPLiFT achieves the best performance when multiple depths are used during training. For our primary results, we use $d \in \mathbf{D} := \{1, 2, 3\}$ for each training step. We present ablations with additional configurations in Appendix C. Overall, our final UPLiFT learning objective can be denoted as:

$$\mathcal{L}_{\text{UPLiFT}} = \sum_{d \in \mathbf{D}} \mathcal{L}_d \quad (5)$$

4. UPLiFT for Predictive Tasks

Experimental Methods. We start with experiments focused on dense predictive tasks: semantic segmentation and monocular depth estimation. We compare with other

task-agnostic feature-upsampling methods LiFT [43], Featup [15], LoftUp [18], JAFAR [10], and AnyUp [51]. We also compare with simple baselines of bilinear and nearest-neighbor upsampling. To compare with prior works, we select DINOv2-S/14 [34] as our primary standard backbone for experiments. For all tasks in this section, we use an UPLiFT model trained for one epoch on the ImageNet-1K dataset [12], with a maximum ground truth image size of 448 and a maximum input image size of 224. We train using our multi-step loss (Section 3.3) with 3 depth levels.

We build on the experimental protocols of [10] and train linear probes models on top of the upsampled DINOv2-S/14 features produced by each method. In addition, following the example of [51], we make a small correction to the learning rate schedule of the official JAFAR evaluation suite. For this reason, we have recomputed all the baseline methods to provide a fair comparison. All methods are run in a pixel-dense upsampling mode. Following the example of [10], we also run LiFT in a $2\times$ upsampling mode followed by bilinear upsampling to pixel-dense features, which we denote as LiFT-2 \times . All evaluations use 448×448 input images and 448×448 upsampled features. For each method, we also report the number of parameters and the average upsampling time for a single image’s feature as measured on one NVIDIA A5000 GPU. Please see the Appendix A for additional details.

Segmentation. We present semantic segmentation results in Table 1 for four datasets: COCO-Stuff [26], Pascal VOC [14], ADE20k [54], and Cityscapes [9]. First, our results confirm the observations of [10], that iteratively running LiFT to produce pixel-dense features leads to degraded performance in this task. We find that a single LiFT upsampling iteration (LiFT-2 \times) yields better performance than iterative upsampling to pixel-dense features. Next, we see that UPLiFT achieves state-of-the-art performance, with the highest mIoU and Accuracy scores for all four datasets,

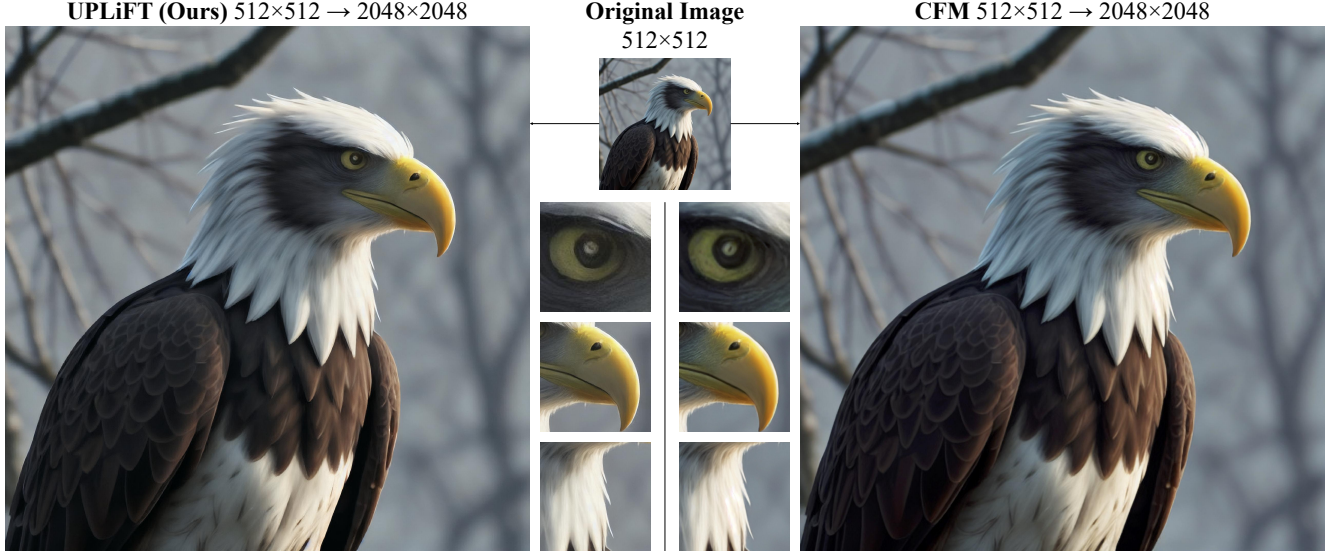


Figure 6. **2048×2048 Image Generation Comparison with CFM.** We show that UPLiFT achieves comparable visual upsampling quality to CFM [15], while using only $1/6^{\text{th}}$ the network parameters, $1/200^{\text{th}}$ the training data, and only 2 iterative upsampling steps.

Table 2. **Zero-shot COCO-5k 512→1024.** UPLiFT achieves superior FID to CFM for COCO text-to-image up-sampling, while running with significantly faster inference speeds. Like [41], we measure speed for this task with batch size 4. CFM is run with 20 steps. [†]Row adapted from [41]

Diffusion Model		Upsampler			Performance	
Model	Steps	Method	Params (M)	Resolution	s/img ↓	FID ↓
SD1.5	25	—	—	512	3.96	24.88
SD1.5	25	CFM-20 [†]	306	512 → 1024	8.79	28.81
SD1.5	25	UPLiFT	53	512 → 1024	5.15	24.23
SDXL	50	—	—	1024	39.84	24.07

even surpassing recent methods that use cross-attention feature pooling. This result demonstrates that our Local Attender approach is an effective method for feature pooling and upsampling. Finally, we see that UPLiFT has faster inference speed than all high-performing pixel-dense upsamplers. While other baselines like Nearest, Bilinear, LiFT-2×, and LiFT achieve faster speeds, they also have inferior performance, meaning UPLiFT provides an optimal combination of speed and feature quality.

Depth Estimation. We next apply UPLiFT to monocular depth estimation on COCO-Stuff. Like [10], we report the δ_1 score, which is a thresholded accuracy score, and the Root Mean Square Error (RMSE). Predicting visual depth from monocular cues requires a broader understanding of the entire image. In theory, the cross-attention-based methods, LoftUp, JAFAR, and AnyUp, should have an advantage in this task, as their feature pooling approach can gather broader information from wider image regions. However, we find that UPLiFT ties with AnyUp for the lowest RMSE score, and is second only to AnyUp for δ_1 score. These

Table 3. **Zero-shot reLAION-400M-5k 512→1024.** UPLiFT demonstrates strong upsampling for text-to-image, surpassing CFM in all three metrics, with faster inference. CFM is run with 4 or 40 steps. *As CFM does not report diffusion steps, we include two options for comparison. Speed measured with batch size 1 as in [41]. [†]Rows adapted from [41].

Diffusion Model		Upsampler			Performance			
Model	Steps	Method	Params (M)	Resolution	s/img ↓	CLIP ↑	FID ↓	p-FID ↓
SD1.5	25	—	—	512	1.27	31.67	22.32	—
SD1.5	40	—	—	512	1.96	31.87	21.90	—
SD1.5	*	CFM-4 [†]	306	512 → 1024	2.72	26.16	21.61	15.83
SD1.5	*	CFM-40 [†]	306	512 → 1024	3.16	26.14	21.67	15.96
SD1.5	25	UPLiFT	53	512 → 1024	1.56	31.00	21.17	16.10
SD1.5	40	UPLiFT	53	512 → 1024	<u>2.27</u>	31.17	20.73	15.49
LCM SDXL	4	—	—	1024	0.72	31.06	28.05	24.13

results demonstrate that our UPLiFT features, which are upsampled using only local information through our Local Attender, can still derive sufficient global information from the backbone features to achieve competitive performance in depth estimation.

Efficiency and Scaling of UPLiFT. As shown by Table 1, UPLiFT achieves faster inference speeds than recent state-of-the-art methods using cross-attention pooling. Note that those speeds are measured for an image size of 448×448 with a patch size of 14, which yields 1024 visual tokens. As recent methods have moved to using cross-attention for feature scaling, they now face quadratic time and memory scaling with the number of input tokens. Meanwhile, our UPLiFT method maintains linear scaling, leading to even faster performance for larger image sizes. This is essential, as producing high-resolution image features can easily become a memory-intensive process. To demonstrate this, we test UPLiFT, LoftUp, JAFAR, and AnyUp for gradually increasing image sizes, until they hit a memory limit of 24 GB on an A5000 GPU. We assume a patch size of 16 and run

Table 4. **Zero-shot Image Super Resolution 256→1024 on 5k samples of FacesHQ and LHQ.** UPLiFT achieves superior scores in SSIM and PSNR for LHQ, and better SSIM for FacesHQ. We note that UPLiFT achieves these results with a single general-purpose module, while CFM uses dataset-specific trained modules with significantly more parameters and iterations. [†]Row adapted from [41]

Upsampler		FacesHQ					LHQ				
Method	Iters	Params (M)	SSIM ↑	PSNR ↑	FID ↓	p-FID ↓	Params (M)	SSIM ↑	PSNR ↑	FID ↓	p-FID ↓
Bilinear	1	–	0.72	21.74	108.12	176.90	–	0.63	22.09	147.65	224.70
Nearest	1	–	0.66	20.48	193.08	332.28	–	0.60	21.08	241.55	288.66
CFM [†]	50	113	<u>0.82</u>	30.40	1.36	1.62	306	<u>0.69</u>	<u>25.69</u>	2.27	2.38
UPLiFT	2	53	0.84	<u>29.93</u>	<u>5.20</u>	<u>6.62</u>	53	0.73	26.70	<u>5.29</u>	<u>7.23</u>

all methods in a pixel-dense upsampling configuration. We report the average inference speed against the number of visual tokens in Figure 1. We find that all baseline methods run out of memory at roughly 1500 visual tokens, which corresponds to an image size of 624×624. At this image size, UPLiFT has inference speeds 2.5-5× faster than the baseline methods. UPLiFT can upsample images with up to 2601 visual tokens without hitting our memory limit.

5. UPLiFT for Generative Tasks

Experimental Methods. In this section, we extend UPLiFT to generative tasks by applying it to VAE latent features, enabling efficient image generation and super-resolution. To preserve the feature distribution for generative tasks, we find that a larger UPLiFT model is a necessity, but our larger model still only has one half or one sixth the parameters of the comparable CFM models. We train UPLiFT for 5 epochs on Unsplash-Lite [46], which contains 25k high-quality images. For comparison, CFM’s general-purpose model is trained on the Unsplash dataset with over 5 Million images. Using a maximum ground-truth image size of 1024, we train across 4 depth levels using the multi-step iterative loss introduced in Section 3.3. Both methods use the Stable Diffusion 1.5 [38] VAE to perform encoding and decoding between pixel and latent space.

We follow the protocols of [41] across two generative tasks: text-to-image diffusion upscaling and image super-resolution. We first use 5k random samples from COCO 2014 and reLAION-400M to evaluate UPLiFT’s ability to upscale diffusion features. This task effectively increases Stable Diffusion 1.5’s output size from 512×512 to 1024×1024. We compare with CFM, as well as SDXL and LCM-LoRA SDXL, which natively generate at 1024×1024, to provide an additional point of reference for visual quality from high-resolution latents. We also evaluate UPLiFT for image super-resolution on the FacesHQ and LHQ datasets, again following the protocols of [41]. Additionally, we compare against two lightweight baselines: direct bilinear and nearest neighbors upsampling in latent space. Reported latency is measured on a single NVIDIA A100 GPU. Please see Appendix B for addition details on the model, training, and tasks.

Efficient High-Resolution Image Generation. On COCO (Tab. 2), UPLiFT achieves a lower FID than CFM in a single step while reducing latency by 41%. UPLiFT is also close to matching the FID of SDXL, which natively generatives 1024 images at much slower speed. On reLAION (Tab. 3), UPLiFT achieves strong performance and faster speeds than CFM, with our 40-step configuration producing the highest CLIP, best FID and patch-FID, while maintaining significantly lower latency. These results demonstrate UPLiFT as a strong choice for efficient generative feature upscaling.

Image Super-Resolution. We evaluate UPLiFT for 4× image super-resolution via latent-space-upsampling, comparing against two lightweight baselines: bilinear and nearest neighbors. Note that the baseline CFM models for this task are trained specifically for the datasets tested, while UPLiFT uses the same general-purpose model from the previous section. Despite this, UPLiFT achieves competitive visual quality using just two iterative steps, with an inference time of only 271 milliseconds per image. Remarkably, UPLiFT is just 13% slower than bilinear upsampling in latent space, yet delivers an order-of-magnitude improvement in super-resolution quality. These results highlight UPLiFT’s efficiency and generalization capabilities.

6. Conclusion

In this work, we have presented UPLiFT, an efficient feature-upsampling method to create high-resolution feature maps from low-resolution features of pre-trained visual backbones. We show that UPLiFT produces high-performing pixel-dense features, and it does so with lower inference costs than existing comparable methods. UPLiFT achieves this through an iterative, convolutional upsampling architecture, without need for expensive cross-attention. Instead, we propose an efficient Local Attender operator, which succeeds in maintaining semantic consistency of features for minimal extra computational cost. Finally, we show that UPLiFT achieves state-of-the-art performance in a range of tasks, including both predictive and generative domains. Our code and UPLiFT models will be released at time of publication. We hope that UPLiFT will help to improve the efficiency and practicality of deep models for dense visual tasks.

References

- [1] Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023. 3
- [2] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021. 1
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 13
- [4] 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
- [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 2
- [6] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9640–9649, 2021. 2
- [7] Yu Chen, Ying Tai, Xiaoming Liu, Chunhua Shen, and Jian Yang. Fsrnet: End-to-end learning face super-resolution with facial priors. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2492–2501, 2018. 3
- [8] Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8628–8638, 2021. 2
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 6
- [10] Paul Couairon, Loïck Chambon, Louis Serrano, Jean-Emmanuel Haugeard, Matthieu Cord, and Nicolas Thome. Jafar: Jack up any feature at any resolution. *arXiv preprint arXiv:2506.11136*, 2025. 1, 2, 3, 4, 6, 7, 12, 16
- [11] Ryan Dahl, Mohammad Norouzi, and Jonathon Shlens. Pixel recursive super resolution. In *Proceedings of the IEEE international conference on computer vision*, pages 5439–5448, 2017. 3
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [13] Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [14] Mark Everingham, SM Ali Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015. 6
- [15] Stephanie Fu, Mark Hamilton, Laura Brandt, Axel Feldman, Zhoutong Zhang, and William T Freeman. Featup: A model-agnostic framework for features at any resolution. *arXiv preprint arXiv:2403.10516*, 2024. 1, 2, 6, 7, 12, 16, 18
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2
- [17] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. 2
- [18] Haiwen Huang, Anpei Chen, Volodymyr Havrylov, Andreas Geiger, and Dan Zhang. Loftup: Learning a coordinate-based feature upsampler for vision foundation models. *arXiv preprint arXiv:2504.14032*, 2025. 1, 2, 3, 4, 6, 12, 13, 16
- [19] Juno Hwang, Yong-Hyun Park, and Junghyo Jo. Upsample guidance: Scale up diffusion models without training. *arXiv preprint arXiv:2404.01709*, 2024. 3
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. pmlr, 2015. 13
- [21] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2, 3
- [22] Johannes Kopf, Michael F Cohen, Dani Lischinski, and Matt Uyttendaele. Joint bilateral upsampling. *ACM Transactions on Graphics (ToG)*, 26(3):96–es, 2007. 2, 12
- [23] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, Sumith Kulal, Kyle Lacey, Yam Levi, Cheng Li, Dominik Lorenz, Jonas Müller, Dustin Podell, Robin Rombach, Harry Saini, Axel Sauer, and Luke Smith. Flux.1 kontext: Flow matching for in-context image generation and editing in latent space, 2025. 3
- [24] 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 3
- [25] Yuming Li, Peidong Jia, Daiwei Hong, Yueru Jia, Qi She, Rui Zhao, Ming Lu, and Shanghang Zhang. Asgdifusion: Parallel high-resolution generation with asynchronous structure guidance. *arXiv preprint arXiv:2412.06163*, 2024. 3
- [26] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6
- [27] Zhihang Lin, Mingbao Lin, Wengyi Zhan, and Rongrong Ji. Accdiffusion v2: Towards more accurate higher-resolution

- diffusion extrapolation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 3
- [28] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022. 3
- [29] Wenze Liu, Hao Lu, Hongtao Fu, and Zhiguo Cao. Learning to upsample by learning to sample. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6027–6037, 2023. 2
- [30] Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*, 2022. 3
- [31] Hao Lu, Wenze Liu, Zixuan Ye, Hongtao Fu, Yuliang Liu, and Zhiguo Cao. Sapa: Similarity-aware point affiliation for feature upsampling. *Advances in Neural Information Processing Systems*, 35:20889–20901, 2022. 2
- [32] Simian Luo, Yiqin Tan, Suraj Patil, Daniel Gu, Patrick Von Platen, Apolinário Passos, Longbo Huang, Jian Li, and Hang Zhao. Lcm-lora: A universal stable-diffusion acceleration module. *arXiv preprint arXiv:2311.05556*, 2023. 14
- [33] Sachit Menon, Alexandru Damian, Shijia Hu, Nikhil Ravi, and Cynthia Rudin. Pulse: Self-supervised photo upsampling via latent space exploration of generative models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2437–2445, 2020. 3
- [34] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 1, 2, 6, 12, 13, 16
- [35] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint arXiv:2307.01952*, 2023. 14
- [36] Haonan Qiu, Shiwei Zhang, Yujie Wei, Ruihang Chu, Hangjie Yuan, Xiang Wang, Yingya Zhang, and Ziwei Liu. Freescale: Unleashing the resolution of diffusion models via tuning-free scale fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16893–16903, 2025. 3
- [37] 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. 2
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 3, 8, 13
- [39] 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. 3
- [40] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Image super-resolution via iterative refinement. *IEEE transactions on pattern analysis and machine intelligence*, 45(4):4713–4726, 2022. 3
- [41] Johannes Schusterbauer, Ming Gui, Pingchuan Ma, Nick Stracke, Stefan Andreas Baumann, Vincent Tao Hu, and Björn Ommer. Fmboost: Boosting latent diffusion with flow matching. In *European Conference on Computer Vision*, pages 338–355. Springer, 2024. 2, 3, 7, 8, 13, 14, 15, 23
- [42] Oriane Siméoni, Huy V Vo, Maximilian Seitzer, Federico Baldassarre, Maxime Oquab, Cijo Jose, Vasil Khalidov, Marc Szafraniec, Seungeun Yi, Michaël Ramamonjisoa, et al. Dinov3. *arXiv preprint arXiv:2508.10104*, 2025. 1, 2, 12, 16
- [43] Saksham Suri, Matthew Walmer, Kamal Gupta, and Abhinav Shrivastava. Lift: A surprisingly simple lightweight feature transform for dense vit descriptors. In *European Conference on Computer Vision*, pages 110–128. Springer, 2024. 1, 2, 3, 6, 12, 16
- [44] Athanasios Tragakis, Marco Aversa, Chaitanya Kaul, Roderick Murray-Smith, and Daniele Faccio. Is one gpu enough? pushing image generation at higher-resolutions with foundation models. *arXiv preprint arXiv:2406.07251*, 2(3):5, 2024. 3
- [45] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv preprint arXiv:2502.14786*, 2025. 2
- [46] Unsplash. Unsplash Full, Lite Dataset 1.3.0, 2025. Accessed: 14 November 2025. 8
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [48] Matthew Walmer, Saksham Suri, Kamal Gupta, and Abhinav Shrivastava. Teaching matters: Investigating the role of supervision in vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7486–7496, 2023. 2, 4
- [49] Jiaqi Wang, Kai Chen, Rui Xu, Ziwei Liu, Chen Change Loy, and Dahua Lin. Carafe: Content-aware reassembly of features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3007–3016, 2019. 2
- [50] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004. 14
- [51] Thomas Wimmer, Prune Truong, Marie-Julie Rakotosaona, Michael Oechsle, Federico Tombari, Bernt Schiele, and Jan Eric Lenssen. Anyup: Universal feature upsampling. *arXiv preprint arXiv:2510.12764*, 2025. 1, 2, 3, 4, 6, 12, 16
- [52] Zhen Yang, Guibao Shen, Liang Hou, Mushui Liu, Luozhou Wang, Xin Tao, Pengfei Wan, Di Zhang, and Ying-Cong

- Chen. Rectifiedhr: Enable efficient high-resolution image generation via energy rectification. *arXiv e-prints*, pages arXiv-2503, 2025. [3](#)
- [53] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023. [2](#)
- [54] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *International Journal of Computer Vision*, 127(3):302–321, 2019. [6](#)
- [55] Minghao Zhou, Hong Wang, Yefeng Zheng, and Deyu Meng. A refreshed similarity-based upsampler for direct high-ratio feature upsampling. *arXiv preprint arXiv:2407.02283*, 2024. [2](#)

A. Additional Details for Predictive Tasks

In this work, we focus on comparing UPLiFT with other task-agnostic feature upsamplers, which have grown in popularity in recent years. Such methods take existing pre-trained feature extractors, and provide an upsampler add-on to provide dense, powerful features directly out of the box. We compare with the following methods:

FeatUp [15]: We compare with the JBU FeatUp variant, which is an iterative upsampler that performs $16\times$ upsampling with a stack of four modified Joint Bilinear Upsamplers [22]. While the implicit variant of FeatUp produces impressive results, it has been shown by works like [18] that this approach is too slow to be practical for large-scale evaluations, with functional inference speeds over $500\times$ slower than comparable methods. We use the official FeatUp JBU model distributed by [15] trained for DINOv2-S/14 without backbone normalization.

LiFT [43]: LiFT is also an iterative upsampling method, which in its base form performs $2\times$ upsampling. The same LiFT model can also be applied iteratively four times to perform $16\times$ upsampling. However, [10] has shown that this iterative upsampling can lead to semantic drift and degraded features. For this reason, we follow the example of [10] and present LiFT in two configurations: “LiFT” runs the model four times for $16\times$ upsampling and “LiFT-2 \times ” runs the model only once, then follows this with bilinear upsampling to pixel scale. As there is not an official LiFT model for DINOv2-S/14, we train one for this work.

LoftUp [18]: LoftUp uses cross-attention to directly up-sample low-resolution backbone features to pixel-scale features, with a $14\times$ upsampling step. We compare with the official LoftUp for DINOv2-S/14 model distributed by [18]. We note that this LoftUp model was originally trained with the DINOv2-S/14 backbone distributed by [34], while other recent methods (ours included) use the version distributed by Hugging Face¹. While these DINOv2 models are fundamentally the same model trained on the LVD-142M dataset and should produce similar results, we choose to conduct an additional test to determine if this has any impact on the performance of LoftUp. We perform evaluation twice with LoftUp, once with each version of DINOv2-S/14, and report the results in Table 5. We find that LoftUp’s performance is nearly identical for both backbones, with only minor variations seen in the mIoU scores for VOC and Cityscapes. For this reason, we conclude that these minor variations in backbone distributions are not a significant confounding factor in our results, and we thus choose to present results for LoftUp with its original backbone for sake of fair representation.

¹https://huggingface.co/timm/vit_small_patch14_dinov2.lvd142m

JAFAR [10]: JAFAR also uses a cross-attention-based approach to perform direct $14\times$ feature upsampling. We build on the evaluation protocols of JAFAR for segmentation and depth estimation, and we compare with the official JAFAR for DINOv2-S/14 model from [10].

AnyUp [51]: AnyUp proposes a modified version of the JAFAR architecture which is designed with an additional “feature-agnostic layer” that enables the model to generalize to different backbones at inference time. Note that AnyUp is still initially trained with features from a particular backbone for upsampling training. Moreover, the results of [51] show that performing inference with a different backbone than was used in training leads to poorer downstream performance than using the same backbone for both training and evaluation, suggesting that training model-specific upsamplers is still preferable for performance. For this study, we compare with the official AnyUp model distributed by [51].

Upsampling Rates. For our main evaluation, we use DINOv2-S/14 as the standard backbone, and run all methods in a configuration to produce pixel-dense features, which requires $14\times$ upsampling. For methods that perform $16\times$ upsampling, which includes our UPLiFT, we follow the example of [10] and simply over-upsample the features, and then downsample to pixel-resolution. We note that this puts UPLiFT at a computational efficiency disadvantage compared to cross-attention-based methods, but despite this, we still achieve faster inference than the cross-attention baselines. Moreover, we note that more recent self-supervised backbones, like DINOv3 [42], have moved back to using patch size 16, which makes $16\times$ upsampling the likely desired option for future research.

B. Additional Details for Generative Tasks

B.1. UPLiFT for VAE Features

We describe additional details of our UPLiFT model designed for VAE feature upsampling and generative tasks.

UPLiFT size for VAE. We empirically find that small parameter count upsampling models, like those demonstrated to be effective for predictive downstream tasks, are insufficient for generative downstream tasks. We illustrate this in Figure 7, where we present a comparison with a small UPLiFT model, which has roughly $2.8M$ parameters and is trained using the same training protocols as our main model. In this comparison, we see that the resulting model produces blurry, low quality upsampling, and is not able to capture high-frequency information. We theorize that the necessary model capacity to train an effective feature upsampler for generative tasks is intrinsically larger than the necessary size for predictive tasks. We base this theory on the intuition that predictive tasks are about narrowing information down

Table 5. **Evaluating the impact of minor backbone variations.** We present a comparison of baseline LoftUp [18] running either with its original DINOv2-S/14 backbone from [34], or with an alternate distribution of DINOv2-S/14 provided by Hugging Face, which is commonly used in other upsampling works. Note that both models are theoretically the same model and only represent different distribution platforms. We find that there is minimal variation in the resulting performance, with the largest changes being in VOC mIoU and Cityscapes mIoU. Our UPLiFT surpasses LoftUp’s performance for either backbone. We report results for LoftUp with its original backbone in the main work to provide a fair representation of the method.

Upsampler			Semantic Segmentation								Depth Estimation	
			COCO		VOC		ADE20K		Cityscapes		COCO	
Method	Params (M)	Time (ms)	mIoU ↑	Acc ↑	mIoU ↑	Acc ↑	mIoU ↑	Acc ↑	mIoU ↑	Acc ↑	δ_1 ↑	RMSE ↓
LoftUp + Orig Backbone	4.3	223.5	62.19	81.35	84.63	96.33	42.16	75.79	62.09	93.55	58.69	0.68
LoftUp + HF Backbone	4.3	223.5	62.20	81.35	84.52	96.30	42.16	75.77	62.17	93.56	58.70	0.68
UPLiFT	0.8	79.4	62.55	81.57	85.21	96.51	42.97	76.00	65.38	94.41	61.16	0.63

into a compressed understanding, which can be represented more compactly with a smaller network, while generative tasks require a high level of detail over a diverse range of visual textures and patterns to achieve high quality. For this reason, we train a larger UPLiFT model for generative tasks, increasing the number of layers in both the encoder and decoder module and increasing the channels per layer. This larger size UPLiFT has 53.5M parameters, which still makes it significantly smaller than the compared CFM [41] models, which have 113M or 306M parameters.

Refiner Block. We find that it is beneficial to introduce an additional “Refiner Block” after the Local Attender module, which is designed to realign the output features with the distribution expected by the VAE decoder. We demonstrate the importance of the refiner block by ablating it in Figure 7 (Right). Without the refiner block, the upsampled images have significant blocky artifacts, which are likely the result of the strict, linear-combination feature upsampling approach. We note that our Local Attender module, and the cross-attention modules used by recent works, act as a form of strong feature regularization, which ensures that the features output by the upsampler maintain a similar distribution to the original input distribution, which, in theory should also match the distribution expected by the VAE decoder. However, through testing, we find that this design may be too restrictive for a generative context. For this reason, we introduce a post-attender “Refiner Block” which gives UPLiFT an opportunity to improve the final features.

Noise Layers and Augmentation. We follow the example of [41] and concatenate additional gaussian random noise channel inputs at several points in the UPLiFT Encoder module when upsampling VAE features. These additional noise channels are provided to help seed the generation of high resolution details. We also apply gaussian noise augmentation to the input latent representation, following [41].

Color Correction. When training UPLiFT for VAE, all our loss terms are computed in the latent feature space, which improves the efficiency of training by removing the need

to decode high-resolution images. Our UPLiFT training is effective at minimizing the L_2 distance of upsampled features in latent space, however, we find that small perturbations in latent space can lead to slight color shifts after decoding. While this could likely be addressed through the use of pixel-space loss terms, this would greatly increase training costs. Instead, we introduce a simple color correction module after the VAE Decoder at inference time. This module computes the per-color-channel means for the low-resolution input image and the high-resolution output image, and subtracts the difference vector from the output image to re-align the color mean. We find that this simple module is sufficient to remove any minor color shifts.

Layer Normalization. Finally, for our larger VAE UPLiFT model, we find it is beneficial to replace the Batch Norm [20] layers with Layer Norm [3] instead, as it provides better numerical stability for the deeper architecture. When using Batch Norm, we sometimes observe color blob artifacts in the upsampled images, which are a consequence of numerical instability in the model. After replacing the Batch Norm operations with Layer Norm, these artifacts no longer occur.

B.2. Experimental Methods for Generative Tasks

We summarize additional key details of our experimental protocols for generative tasks.

Diffusion Upsampling. We evaluate UPLiFT’s ability to upscale diffusion features on COCO 2014 and reLAION-400M following the protocols of [41]. On COCO, we randomly sample 5k caption-image pairs. On reLAION, we randomly sample 5k images with a minimum resolution of 1024×1024 . While [41] used the now deprecated LAION dataset, our sample is sufficiently similar, and our computed FID with LCM-LoRA SDXL matches that of CFM. During inference, we generate latents with Stable Diffusion 1.5 [38] for each sampled caption, and then decode them to generate 5k 512×512 images. We apply UPLiFT to produce 1024×1024 images. Note that UPLiFT leverages the image created from the decoded low-resolution latents to guide



Figure 7. **Visual ablation of design choices for VAE UPLiFT.** (Left) UPLiFT achieves high quality 512→1024 upsampling with a larger parameter count model. (Middle) A smaller-scale UPLiFT has insufficient capacity to upsample all high-frequency information and produces blurry results. (Right) Ablation of the Refiner Block leads to blocky artifacts in upsampled images.

feature upsampling, similar to CFM, which must decode the image as part of the Pixel-Space Upsampling (PSU) method. However, UPLiFT does not require re-encoding a bilinearly upsampled image like CFM does, which gives an additional efficiency gain to UPLiFT.

The default precision of `float32` is used for all diffusion processes, and each model utilizes its default scheduler with the specified number of steps. We also compare with SDXL [35] and LCM-LoRA SDXL [32], which natively generate at 1024×1024 , to provide an additional point of reference for visual quality from high-resolution latents. SDXL can be viewed as an ‘upper-bound’ oracle for native megapixel generation, while LCM-LoRA SDXL provides a point of reference of a low-latency distilled model at megapixel native generation. Please note that SD 1.5 cannot natively generate latents for images at resolutions different from 512×512 without significant quality degradation or alterations. [41] includes experiments upscaling 256×256 to 1024×1024 by specially fine-tuning a Stable Diffusion 1.5 model at a lower resolution. For our tests, we choose to prioritize experiments using only the official Stable Diffusion 1.5 model generating latents for 512×512 images. Utilizing the official code of [41], we compute FID and patch-FID, which splits the image into four random patches sized 512×512 . For CLIP score, we use the Python package `clip_score` with `clip-vit-base-patch32`.

Super-Resolution. Following CFM [41], we perform super-resolution evaluations on two datasets: FacesHQ and LHQ. We randomly sample 5k images from the LHQ dataset and the joint combination of CelebA-HQ and FFHQ datasets, called FacesHQ. For each image, we perform a 1024×1024 center crop, followed by a bilinear downsample to 256×256 . Finally, we apply UPLiFT to recover the

full sized image. To do so, we first use Stable Diffusion 1.5’s VAE to encode the image into latent space. Then we apply the UPLiFT module with 2 iterations, resulting in a $4 \times$ latent, which is decoded using the same VAE’s decoder back into pixel space. Performance is measured using the evaluation scripts of [41], which track SSIM [50], PSNR, FID, and patch-FID.

Latency Timing. To ensure fair comparison with [41], we measure latency on a single NVIDIA A100-SXM4-80GB across both experiments. We use a batch size of 1 for experiments on reLAION, and 4 for experiments on COCO to be consistent with [41]. All experiments utilize the `torch.compile` module, and measurements are conducted with 3 warm-up batches. The final latency is averaged across 10 subsequent batches.

C. Ablations of UPLiFT Design Choices

C.1. Training Depth

As discussed in Section 3.3, we train UPLiFT with a multi-depth, multi-step feature reconstruction training objective, where the level of downsampling and the number of up-sampling steps depends on the training depth, d . Here we present an ablation over d , where we train our UPLiFT model with different training depth configurations, including individual depths, or multiple concurrent depths. We use semantic segmentation on COCO and VOC with pixel-dense features as our evaluation tasks for this ablation. The results are summarized in Table 6.

When training with only one depth level (rows 1–4), better performance is achieved with shallower depths. VOC has the best performance with $d=1$, and COCO has slightly better performance with $d=2$. We believe this occurs be-

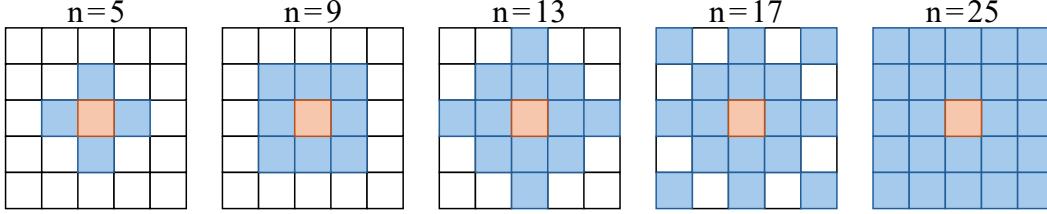


Figure 8. **Local Attender neighborhood designs.** We visualize the neighborhood designs tested with our Local Attender module. The center token (orange) is always included in the neighborhood, and the blue tokens represent the offset relative local positions that are included in feature pooling through local attention.

Table 6. **Ablation of training depth configurations.** We test COCO and VOC segmentation for different training depth levels, including individual and concurrent depth configurations.

Depth	COCO		VOC	
	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow
$d = 1$	62.26	81.34	85.03	96.42
$d = 2$	62.34	81.43	84.89	96.40
$d = 3$	62.24	81.34	84.67	96.33
$d = 4$	60.82	80.30	81.42	95.42
$d \in \{1, 2\}$	62.46	81.50	85.04	96.44
$d \in \{1, 2, 3\}$	62.55	81.57	85.21	96.51
$d \in \{1, 2, 3, 4\}$	62.36	81.45	84.97	96.42

cause shallower depths give the UPLiFT Encoder a large input image, which allows it to learn more about high-frequency image information. For $d=3$, we see a slight drop in performance and for $d=4$ we see a major drop in performance. Under $d=4$, the 448×448 image is downsampled to 24×24 which yields only a 2×2 token grid. It seems this level of downsampling is too severe for effective learning.

Next, we find that training with multiple depths concurrently can enhance performance. As shown by row 5, training with $d \in \{1, 2\}$ concurrently leads to superior performance on both COCO and VOC. This effect is further enhanced by adding in $d=3$. However, the addition of $d=4$ becomes detrimental for multi-depth training, which again indicates that downsampling the input too severely hinders learning. From this, we find that training with a multi-depth configuration with $d \in \{1, 2, 3\}$ yields the best results, so we use this training configuration in our predictive task experiments. Note that for our VAE UPLiFT model, we train with a maximum resolution of 1024×1024 in following [41]. For this larger image size, we find that $d \in \{1, 2, 3, 4\}$ works to good effect.

C.2. Local Attender Neighborhood

We propose a new Local Attender module, which uses a relatively-defined, variably-sized local neighborhood to perform local attention pooling in linear time. In this section, we present additional experiments with different

Table 7. **Ablation of Local Attender neighborhood sizes.** We compare UPLiFT performance with Local Attender modules with different neighborhood patterns, or without the Local Attender.

Neighbors	COCO		VOC	
	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow
No LA	47.52	73.06	64.86	91.37
$n = 5$	62.15	81.30	84.80	84.80
$n = 9$	62.33	81.42	84.99	96.44
$n = 13$	62.51	81.53	85.10	96.47
$n = 17$	62.55	81.57	85.21	96.51
$n = 25$	62.46	81.50	85.22	96.48

neighborhood sizes and shapes. We limit the neighborhood size to a maximum of a 5×5 . We present an illustration of the neighborhood sizes and shapes tested in Figure 8. In addition, we present results for an UPLiFT model with the Local Attender module ablated. All models are trained for one epoch on ImageNet-1k, with 448×448 max image size, like the main work. Results are summarized in Table 7.

First, from row 1, we see that the removal of the Local Attender module significantly harms performance, which indicates that our Local Attender is essential for maintaining effective feature upsampling through multiple steps. Next, in row 2, we test $n=5$ with what we consider to be the smallest viable neighborhood, which consists of only the current token and its immediately touching neighbors. Already for $n=5$ we see strong performance in both datasets, though this performance continues to improve as we expand the neighborhood to 9, 13, and 17 neighbors. The $n=25$ configuration, which compared to $n=17$ adds on several additional neighbors along the second layer of offsets, only leads to roughly equal performance in VOC and slightly lower performance in COCO. This result indicates that the addition of further neighbors may not be beneficial if they are added in intermediate places between existing neighbors. Overall, our best performance is achieved with the star-shaped $n=17$ neighborhood design, so we use this as our Local Attender neighborhood for all our experiments in the main work.

Table 8. **Semantic Segmentation with DINOv3 and UPLiFT**. We measure segmentation performance on COCO and VOC for DINOv2-S/14 vs. DINOv3-S+/16. We compare against JAFAR and AnyUp, which either report results with this DINOv3 backbone or have provided an official model for it [10]. We find UPLiFT gives the best performance for all metrics, datasets, and backbones. [†]Row adapted from [51].

		DINOv2-S/14				DINOv3-S+/16			
Upsampler		COCO		VOC		COCO		VOC	
Method	Params (M)	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow	mIoU \uparrow	Acc \uparrow
JAFAR	0.7	61.71	81.01	<u>84.38</u>	<u>96.22</u>	62.47	81.50	<u>83.05</u>	<u>95.99</u>
AnyUp [†]	0.9	<u>62.16</u>	<u>81.37</u>	84.00	96.19	<u>62.99</u>	<u>81.84</u>	–	–
UPLiFT	0.8	62.55	81.57	85.21	96.51	63.74	82.27	84.72	96.55

D. Additional Results and Visualizations

D.1. UPLiFT for DINOv3

In our primary results in Section 4, we follow the example of recent feature upsampling works and use DINOv2-S/14 [34] as our primary backbone for assessing UPLiFT and comparing it with baseline works. However, as of writing, DINOv3 [42] has recently been published, and we expect future work in feature upsampling will likely shift to focus on this family of backbones. For this reason, we present additional results for Semantic Segmentation using DINOv3-S+/16. We compare with JAFAR [10], which has released an official model for this backbone on their repository, and with AnyUp [51], which has published results for DINOv3 on COCO in their work. While we would additionally like to compare with LoftUp [18], no official DINOv3 LoftUp upsampler has been published as of writing, nor has official training code been made publicly available. We present results for COCO and VOC in Table 8, alongside results for DINOv2-S/14 for comparison. We see that the DINOv3 features lead to stronger performance in COCO, with improved results for DINOv3 with UPLiFT. However, DINOv3 does not necessarily yield stronger performance in VOC, as is seen for both UPLiFT and JAFAR. However, for both the DINOv2 and DINOv3 backbones, UPLiFT achieves the best semantic segmentation performance for both metrics and datasets.

D.2. Additional Efficiency Comparisons

In Section 4, we assessed the efficiency of UPLiFT in comparison to recent state-of-the-art cross-attention-based feature upsamplers [10, 18, 51]. We present additional efficiency results here. Along with reporting inference time, we also report the inference GPU memory usage, and we also present these results for two different GPU types: an NVIDIA A5000 with 24 GB of memory, and an NVIDIA A6000 with 48 GB of memory. We plot these results in Figure 9. First, we see that the cross-attention-based methods clearly show quadratic scaling in their memory usage. All three methods reach the A5000’s 24 GB memory limit at around 1500 visual tokens, while UPLiFT can reach over

2500 tokens. We note that at the lower range of image sizes, UPLiFT uses slightly more memory than the baseline methods, but by 1000 tokens, UPLiFT’s linear scaling makes it more efficient. Moreover, the improved efficiency of UPLiFT is even more apparent when tested on an A6000. With twice the GPU memory, UPLiFT can process twice as many visual tokens on an A6000, while the other methods rapidly fill the extra memory. The improved speed of UPLiFT for larger images is also apparent on the A6000, with it being 2–4 \times faster than the baseline methods at 2000 tokens.

D.3. Semantic Stability with UPLiFT

A critical issue with LiFT [43] is the occurrence of semantic drift through iterative feature upsampling. We present a visual comparison of semantic drift in LiFT compared with the semantic stability achieved by UPLiFT. Following [15], we visualize the features of both methods using Principle Component Analysis (PCA) with 3 components. We extract the backbone features and all intermediate feature upsampling steps from both methods, and then perform a joint-PCA over the full collection. This means that consistent colors from one image to the next means consistent features too. We visualize the results in Figure 10.

Examining the PCA images for LiFT (row 1), we see that iterative upsampling leads to increasing degradation of the latents. The features become faded, murkier, and distorted through each step. Local image regions with consistent semantic content, like the dog’s ear and nose region (row 3), do not maintain consistent and similar features. In comparison, UPLiFT (row 2) maintains consistent features through each upsampling step, without signs of internal distortions. In addition, the edges of objects are more distinctly defined than LiFT, and they match well to the original image, which enables better performance in tasks like semantic segmentation. Overall, these results demonstrate that our UPLiFT approach, through the use of Local Attender, is able to maintain consistent feature semantics through iterative upsampling steps to produce high-quality, pixel-dense features.

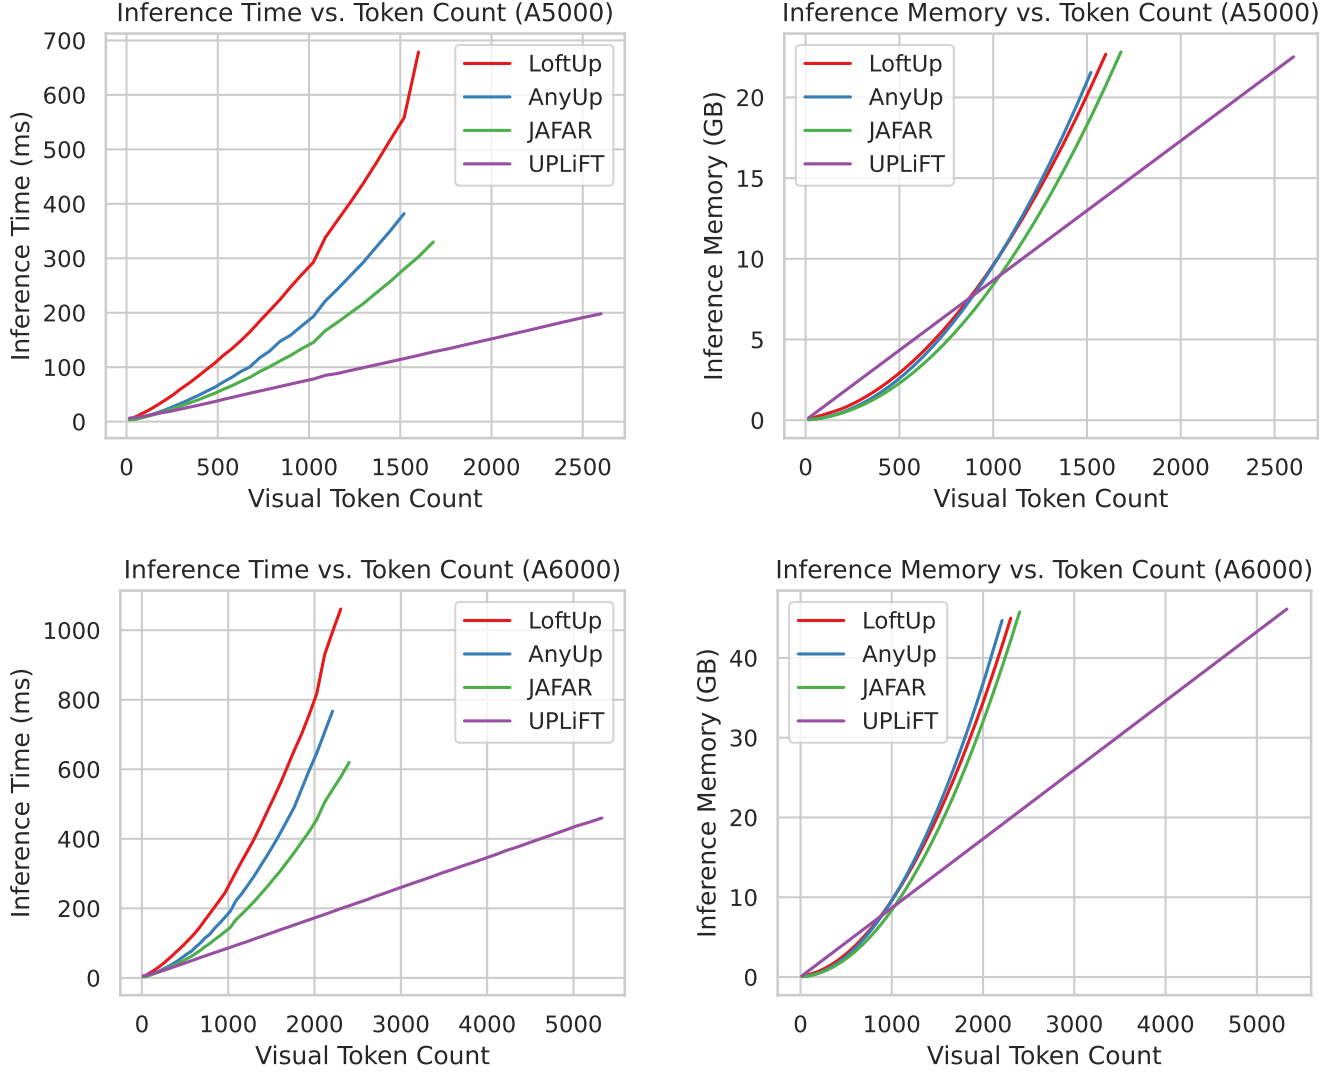


Figure 9. **Speed and Memory Usage Comparison.** We compare UPLiFT with recent cross-attention-based feature upsampling methods when running with gradually increasing image sizes. We report the average inference time and memory usage against the visual token count for two different GPU types: an NVIDIA A5000 with 24 GB of memory and an NVIDIA A6000 with 48 GB. UPLiFT maintains linear time and memory scaling with respect to the number of tokens, while the baseline methods show quadratic scaling. This gives UPLiFT far faster performance as the token count increases, and allows us to run pixel-dense feature upsampling on larger images.

D.4. Additional Visualization for Generative Tasks

We provide further visualizations for UPLiFT applied to generative tasks. Figure 11 highlights UPLiFT’s strong and efficient performance in $4\times$ super-resolution from 512×512 to 2048×2048 . We find that UPLiFT adds only 8.47% latency compared to bilinear upsampling in latent space, and yields far superior image quality to the naive upsampling method. Figure 12 includes 2048×2048 images, up-scaled from Stable Diffusion 1.5, and Figure 13 shows the same at 1024×1024 . For the super-resolution task, we include samples from the evaluation datasets FacesHQ and LHQ using $2\times$ iterative upsampling twice, from 256×256 to 1024×1024 in Figure 14 and Figure 15 respectively.

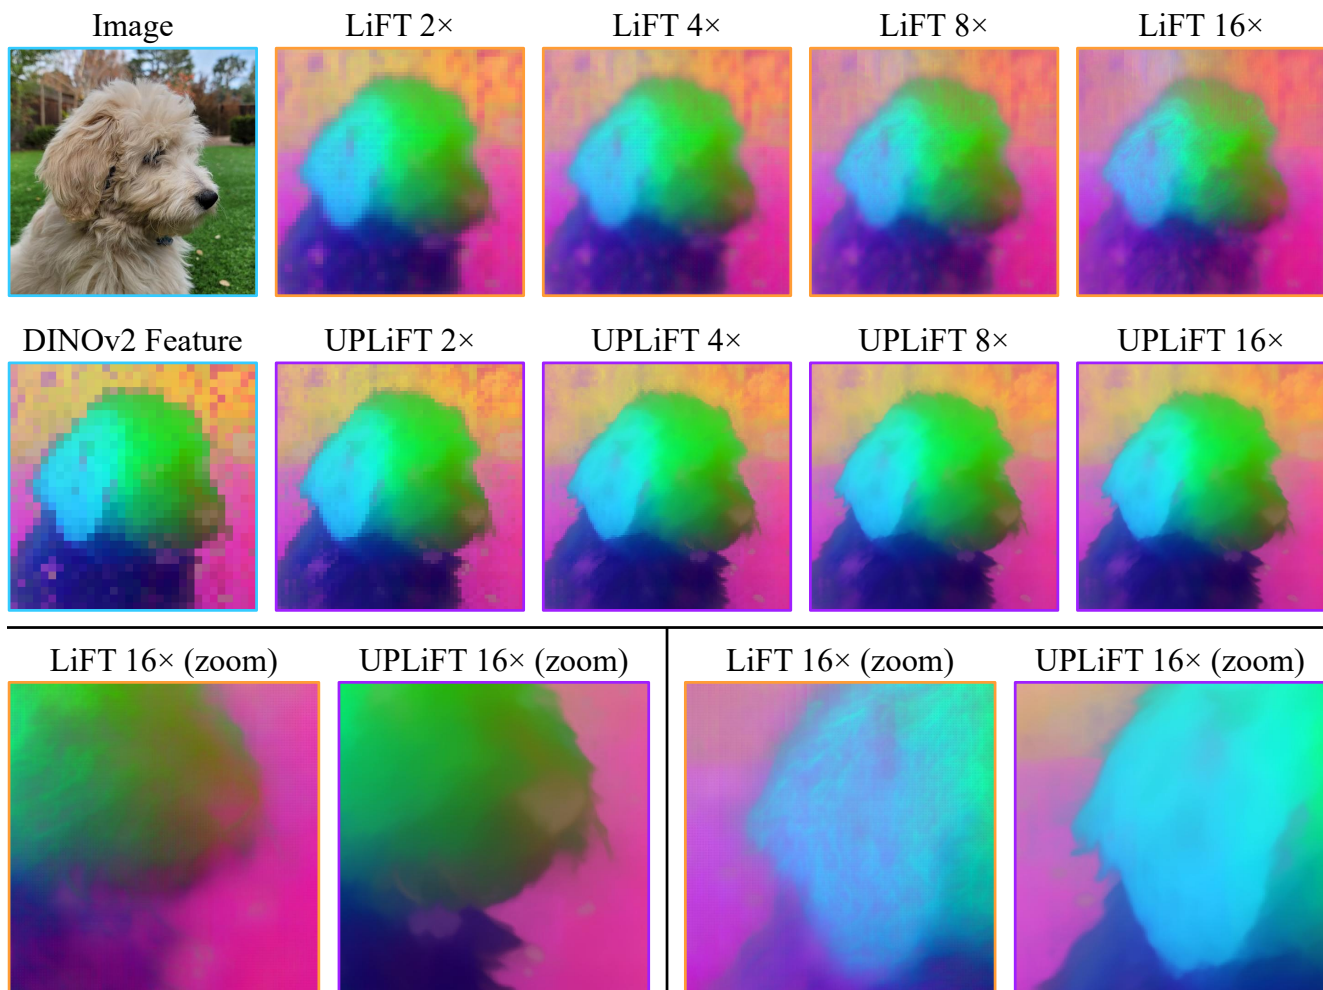


Figure 10. **Comparison of semantic drift in LiFT and semantic stability in UPLiFT.** We visualize intermediate feature upsampling steps through PCA, following [15]. LiFT shows signs of feature drift, with local features becoming murkier and more distorted in deeper steps. This drift can lead to poor performance in downstream tasks, as the strength of the original backbone representation is lost. UPLiFT maintains consistent feature semantics thanks to our Local Attender, and local object regions maintain consistent features (coloration) across all upsampling stages.




Figure 11. **UPLiFT vs. Latent Space Bilinear $4\times$ upsampling for image super-resolution**. We compare the latency of UPLiFT versus applying bilinear upsampling in latent space for image super-resolution to demonstrate UPLiFT’s state-of-the-art efficiency. While only 8.47% slower end-to-end, UPLiFT produces significantly better visual fidelity, as shown by the zoomed-in views (bottom) which display the source low-resolution image compared with UPLiFT’s super-resolution image. The low-resolution image is selected from the FacesHQ dataset and bilinearly downsampled from 1024×1024 to 512×512 before UPLiFT is applied. Best viewed zoomed in.




Figure 12. **UPLiFT $512 \times 512 \rightarrow 2048 \times 2048$ upsampled images using Stable Diffusion 1.5.** We apply our VAE UPLiFT model to this task in a $4\times$ upsampling configuration. UPLiFT upsamples latents corresponding to 512×512 images generated using 50 diffusion steps on Stable Diffusion 1.5. Best viewed zoomed in.




Figure 13. **UPLiFT $512 \times 512 \rightarrow 1024 \times 1024$ upsampled images using Stable Diffusion 1.5.** We apply our VAE UPLiFT model to this task in a $2\times$ upsampling configuration. UPLiFT upsamples latents corresponding to 512×512 images generated using 50 diffusion steps on Stable Diffusion 1.5, and the end-to-end latency is 2.75 seconds on an NVIDIA A100 GPU. The UPLiFT model itself takes only 104 milliseconds of this time. Best viewed zoomed in.



Figure 14. **UPLiFT $256 \times 256 \rightarrow 1024 \times 1024$ super-resolution samples from FacesHQ.** We use our VAE UPLiFT model that is *not fine-tuned* for image super-resolution and is only trained in latent space. Our end-to-end upsampling time is only **270.9 milliseconds** on an NVIDIA A100 GPU. Best viewed zoomed in.




Figure 15. **UPLiFT $256 \times 256 \rightarrow 1024 \times 1024$ super-resolution samples from LHQ.** We use our VAE UPLiFT model, which is trained as a generalist model and is not specifically fine-tuned for this dataset. The LHQ dataset presents a greater challenge than FacesHQ, based on the diversity of visual textures present. Despite this challenge, we see good performance with our generalist UPLiFT. In comparison, [41] uses a fine-tuned model with $3 \times$ the parameter count for evaluations on LHQ versus FacesHQ. Best viewed zoomed in.