

MMDiff: Extending Diffusion Transformers for Multi-Modal Generation

Yagmur Akarken¹, Orest Kupyn¹, and Christian Rupprecht¹

University of Oxford, Visual Geometry Group




Fig. 1: MMDiff. We present a method to decode multiple modalities from a frozen Diffusion Transformer. We show that information in DiT is distributed across the denoising trajectory, and introduce multi-timestep feature fusion to capture the complete perceptual information encoded across them. Lightweight decoder heads predict semantic segmentation, saliency, and depth with no modifications to the generative model.

Abstract. Diffusion transformers have demonstrated remarkable generative capabilities, yet the rich perceptual representations computed across their denoising trajectory are discarded once the content is rendered. We present MMDiff, a framework that transforms a frozen diffusion transformer into a multi-modal generative system that jointly produces images alongside any combination of dense perceptual modalities using lightweight decoder heads. Our central finding is that perceptual information is temporally distributed along the denoising trajectory, and that multi-timestep feature fusion with spatially varying aggregation weights is essential, improving semantic segmentation results by up to 28.7% mIoU over single-timestep extraction. We further adopt concept-driven attention extraction for interpretable spatial guidance, and show that frozen diffusion features are competitive with and complementary to state-of-the-art encoders such as DINOv3. By training only lightweight decoder heads on a frozen backbone, we achieve strong performance in semantic segmentation, salient object detection, and depth estimation, and demonstrate that this framework enables effective synthetic data generation at scale. Code, models, and results are available [here](#).

1 Introduction

Image and video generators are often not used solely to generate visual data but also require additional information about the image, such as segmentation or depth, for use in downstream applications. Diffusion transformers process thousands of tokens across dozens of layers and denoising timesteps to generate a single image or video, yet all intermediate representations are discarded once the final content is rendered. Separate task-specific models are then used to extract semantic, geometric, and structural information from the generated pixels, thereby re-extracting information that may already have been computed during generation. Extending diffusion models to predict new perceptual modalities has clear practical value. A model that generates images while simultaneously predicting their semantic labels, depth maps, or saliency masks can, for example, produce annotated training data at scale, thereby eliminating the need for expensive manual annotation. Similarly, such a pipeline is equally valuable for video and 3D generation, where geometric and semantic understanding must be consistent with the generated content. In both settings, the key challenge is the same: how to best extract perceptual information from a frozen generative backbone. Several works have explored this direction. DatasetDM [46] and Dif-fuMask [47] extract annotations from Stable Diffusion’s U-Net [36] during generation. In parallel, recent methods extend diffusion backbones to jointly predict scene geometry and video [23, 51, 53]. However, all of these approaches extract features from the final denoised latent or the last denoising timestep, discarding all intermediate representations computed during generation. The shift toward transformer-based diffusion models makes this particularly important. State-of-the-art models like FLUX [27], Stable Diffusion 3 [12], and PixArt- α [4] adopt the Diffusion Transformer (DiT) architecture [32]. Unlike U-Nets [37], which organize representations hierarchically through spatial downsampling, DiTs process all tokens at uniform resolution throughout the network [9]. Without spatial hierarchies, perceptual information is spread more uniformly across layers and timesteps rather than concentrated at specific scales, meaning the final denoised latent alone is insufficient for dense prediction.

In this work, we introduce **MMDiff**, a multi-modal diffusion framework that decodes any combination of dense perceptual modalities from a frozen diffusion transformer. Our central finding is that multi-timestep feature fusion is essential: aggregating features across the denoising trajectory yields significantly richer representations than extracting from any single timestep. In our experiments, multi-timestep fusion improves segmentation by +28.7% mIoU over single-timestep extraction. We further leverage concept-driven attention extraction [16] for targeted spatial guidance, and show that discriminative features can optionally complement diffusion representations [38]. By training only lightweight decoder heads from a frozen DiT [27] backbone, we achieve strong performance across semantic segmentation, salient object detection, and depth estimation. As a downstream application, this enables automatic label extraction for synthetic data generation. We summarize our contributions as follows:

- We present **MMDiff**, a framework for extending a frozen DiT to predict any dense perceptual modality using lightweight decoder heads, without modifying the generative backbone.
- We demonstrate that multi-timestep feature fusion is essential for dense prediction from DiTs, and that single-timestep extraction, as used by all prior works, discards highly valuable information encoded during generation. We further enrich the extracted features using concept-driven attention extraction [16], which provides richer spatial guidance than averaged cross-attention.
- We show, for the first time, that frozen diffusion transformer features achieve comparable performance to discriminative representations such as DINO [38], and that the two are complementary, demonstrating an effective combination of generative and discriminative features.
- MMDiff enables effective synthetic data generation: a single frozen model simultaneously generates images and extracts high-quality annotations, substantially outperforming prior U-Net-based methods [31, 46, 47].

2 Related Work

Diffusion Transformer Architectures. Diffusion models [8, 18, 39] have transitioned from U-Net [37] to transformer-based architectures for improved scalability and generation quality. Latent diffusion models [36] enable efficient high-resolution synthesis by operating in a compressed latent space. DiT [32] demonstrated that transformers scale more efficiently than U-Nets for image generation, processing all patches with equal capacity through self-attention rather than hierarchical downsampling. FLUX [27] extends this with a dual-stream architecture: 19 double-stream blocks for text-image interaction followed by 38 single-stream blocks for unified processing. Unlike U-Nets, which create semantic hierarchies through spatial downsampling, FLUX maintains 1024 image tokens of dimension 3072 across all layers. Unlike U-Nets, semantic specialization in DiTs arises from attention patterns across layers and timesteps rather than from spatial hierarchies. This fundamentally changes how perceptual information is organized, requiring new strategies to extract it effectively.

Multi-Model Diffusion. Image and video generators often require additional information about the generated content, such as segmentation or depth, for downstream applications. Recent work extends diffusion models to predict geometry alongside generation, though most approaches train separate decoder networks that operate independently of the generation process. For video, FantasyWorld [7] trains a separate geometric branch with DPT [35] heads to decode depth, point maps, and camera poses alongside a frozen video foundation model, while JOG3R [19] jointly trains a separate DUST3R-style [44] geometry decoder with the video generator. WVD [52] trains the diffusion transformer itself to output geometry as separate output channels alongside RGB, though geometry is only available after the full generation process completes via the VAE decoder.




Fig. 2: MMDiff Pipeline: Given a text prompt, the diffusion transformer generates an image through its dual-stream and single-stream blocks. Features are extracted, pruned, and combined with concept attention maps and discriminative features, then decoded to produce task-specific predictions for the generated image.

For images, DatasetDM [46] trains perception decoders on Stable Diffusion features to generate segmentation masks and depth maps, though the coarse spatial resolution of U-Net feature maps limits annotation quality. DiffuMask [47] is similarly constrained to single-concept segmentation by its reliance on cross-attention maps, while Dataset Diffusion [31] addresses multi-class scenarios but faces the same U-Net feature limitations. GenPercept [48] and Diff-2-in-1 [55] integrate customized perception decoders with U-Net diffusion features for various dense prediction tasks. However, these U-Net-based approaches encode semantic information hierarchically through spatial downsampling, limiting the richness of extracted features. Our work demonstrates that transformer-based diffusion models distribute semantic information across layers and denoising timesteps without spatial downsampling. By recycling features computed during generation and aggregating them across multiple timesteps with learned spatial weighting, we extract substantially higher-quality annotations than U-Net-based methods using lightweight decoders on frozen transformer features.

3 Method

Our goal is to extend a frozen diffusion transformer into a multi-modal generator by extracting and fusing features computed across the denoising process. Given a text prompt, we simultaneously generate an image and extract from the DiT internal representations any combination of dense perceptual modalities, such as semantic segmentation, depth, saliency, and more, using lightweight decoder heads. This transforms a single generative model into a pipeline that jointly generates images with any combination of semantic, geometric, and structural descriptions in a single generation pass.

Modern DiTs, such as FLUX [27], employ a dual-stream architecture comprising double-stream blocks followed by single-stream transformer blocks. These models maintain a flat architecture in which all tokens are processed at a uniform spatial resolution throughout the network. Unlike U-Nets, which create hierarchi-

cal representations through spatial downsampling [54], this flat transformer architecture achieves semantic abstraction through attention mechanisms, thereby necessitating different feature extraction strategies. For example, diffusion transformer blocks produce features with very large activation magnitudes [14], necessitating careful normalization when aggregating features across multiple layers and timesteps to ensure stable training and effective gradient flow.

Our framework serves two purposes: (1) generating synthetic datasets with automatically extracted annotations during image generation, and (2) extracting features from real images for validation and analysis. For real images, we apply the forward diffusion process [41] to add Gaussian noise:

$$x_t = \sqrt{\alpha_t}x_0 + \sqrt{1 - \alpha_t}\epsilon, \quad (1)$$

where x_0 is the clean image, $\epsilon \sim \mathcal{N}(0, I)$ is random noise, and t indexes the diffusion timestep. We then extract features from these noised images at multiple timesteps, enabling the same feature extraction mechanism to work for both synthetic generation and real image analysis.

3.1 Multi-Timestep Feature Extraction

We extract features from multiple layers $L_i \in \mathcal{L}$ within the DiT. We choose a variety of early and late layers to provide coverage across the transformer depth while maintaining computational efficiency.

Multi-Timestep Extraction. Semantic information in diffusion transformers is distributed across the denoising trajectory [22]. Early timesteps capture coarse object layouts and scene-level semantics. Late timesteps preserve fine details and local structure but lose global semantic context [50]. To access the full range of semantic knowledge, we extract features from multiple timesteps $t_j \in \mathcal{T}$ in the denoising schedule, where $t \in [0, 1]$ with $t = 1$ representing pure noise and $t = 0$ representing clean data. Combined with the layer selections, this yields $|\mathcal{L}||\mathcal{T}|$ feature maps total per input. We denote the extracted feature at layer i and timestep t as $F_t^{(i)} \in \mathbb{R}^{B \times N \times C}$, where N is the number of spatial tokens and C is the feature dimension.

Token Aggregation. Since each feature map typically contains upwards of 1,000 tokens, each with several thousand channels, naive concatenation would yield prohibitively large representations, leading to memory overflow and computational infeasibility. Simple averaging, while memory-efficient, fails to leverage the complementary nature of features across timesteps and layers. Figure 3 analyzes DiT feature characteristics to motivate our aggregation strategy. Figure 3a shows that activation magnitudes grow monotonically across FLUX’s 38 single-stream layers with high correlation between adjacent layers, indicating gradual information propagation. This motivates our sparse layer sampling at $\mathcal{L} = \{4, 15, 26, 37\}$ to maximize coverage while avoiding redundancy. Figure 3b visualizes feature entropy across timesteps $\mathcal{T} = \{0.78, 0.52, 0.26, 0.001\}$, demonstrating distinct information content at each denoising stage: early timesteps




Fig. 3: Analysis of DiT features. (a) Activation magnitudes grow monotonically across layers with high inter-layer correlation, motivating sparse layer sampling. (b) Feature entropy varies substantially across timesteps, capturing complementary coarse-to-fine information that justifies learned aggregation.

capture coarse layouts with high entropy, middle timesteps refine object boundaries, and late timesteps preserve fine details. The complementary nature of these features necessitates learned spatially-varying aggregation rather than simple averaging.

The module first applies learned linear projections to reduce the feature dimension, producing $F_t^{\text{proj}} \in \mathbb{R}^{B \times N \times C'}$ where $C' < C$, followed by layer normalization. Since diffusion transformer blocks produce features with sometimes very large activation magnitudes, this normalization is critical for stable training. Additionally, different layers exhibit vastly different ranges, which this normalization mitigates.

The core aggregation mechanism is a 3-layer transformer with 8 attention heads and a hidden dimension of 768. Each layer includes pre-normalization, multi-head self-attention with temporal positional embeddings, and feedforward networks with GELU [17] activation. We use LayerScale initialization [42] (10^{-4}) for stability and dropout [40] (0.1) for regularization. The transformer predicts per-pixel aggregation weights $\alpha \in \mathbb{R}^{|\mathcal{T}|}$ through a specialized output head, enabling each spatial location to adaptively weight the four timesteps based on local content.

The weighted combination is refined using Convolutional Block Attention (CBAM) [45], which applies channel and spatial attention. A residual connection from the cleanest timestep ($t = 0$) preserves fine details. Formally, for each spatial location:

$$F_{\text{fused}} = \text{CBAM} \left(\sum_{t \in \mathcal{T}} \alpha_t \cdot F_t^{\text{proj}} \right) + F_0^{\text{proj}}, \quad (2)$$

where the weights α_t sum to 1 at each location. This enables optimal temporal aggregation to vary spatially—smooth regions benefit from early-timestep consistency, whereas boundaries require late-timestep precision.

3.2 Concept-Driven Attention Extraction

Cross-attention maps between text and image tokens provide spatial guidance for dense prediction. However, averaging attention across all text tokens, including articles, prepositions, and punctuation, dilutes spatial precision [11, 24]. Moreover, in multi-modal DiT architectures, text prompt embeddings are sequentially updated at each layer, making it difficult to extract attention maps for concepts not present in the original prompt without altering the generated image appearance.

We adopt concept-driven attention extraction [16], which addresses these limitations by creating a separate stream of concept embeddings that are not used for image generation. The method injects concept tokens (e.g., “object”, “background”) alongside text and image tokens, with initial embeddings produced by a text encoder. These concept embeddings are updated through a one-directional attention mechanism where concepts attend to both image patches and other concepts, but image tokens do not attend to concepts:

$$o_c = \text{softmax}(q_c k_{xc}^T) v_{xc} \quad (3)$$

where q_c are concept queries, and k_{xc}, v_{xc} are concatenated image-concept keys and values. This one-directional design ensures that concepts gather information from the image without affecting generation.

Spatial maps are then extracted by computing dot-product similarities between concept output vectors o_c and image patch output vectors o_x in the attention output space:

$$\phi(o_x, o_c) = \text{softmax}(o_x o_c^T) \quad (4)$$

This differs from traditional cross-attention maps, which operate in the query-key space. We aggregate information across multiple multi-modal layers by averaging: $\frac{1}{|\mathcal{L}|} \sum_{L=1}^{|\mathcal{L}|} \phi(o_x^L, o_c^L)$ where $|\mathcal{L}|$ denotes the number of multi-modal layers.

This approach provides flexibility to extract task-specific spatial guidance by defining relevant concepts at inference time without modifying the generation process. For semantic segmentation, we use concepts such as “object” and “background” in scene parsing. For depth estimation, “near” and “far” encode geometric relationships. For salient object detection, “salient” and “contour” provide targeted guidance. Importantly, this mechanism eliminates the need to fuse features from external text encoders into diffusion blocks, as done in VPD [54]. This mechanism is unique to multi-modal DiT architectures and cannot be replicated by U-Net-based models.

3.3 Discriminative Feature Complementarity

While our framework relies solely on diffusion features, we find that discriminative representations can optionally complement them. Diffusion features encode semantic relationships from generative training, while discriminative features capture distinct visual characteristics from contrastive learning, and the two are complementary rather than redundant. We optionally integrate DINO-v3 features [38] with our diffusion-based representations. We extract features from DINO-v3-ViT-B/16 and concatenate them with our fused DiT features and concept maps. This creates a comprehensive representation where DiT provides semantic structure, concept maps offer targeted spatial guidance, and DINO-v3 contributes complementary discriminative information. As we demonstrate in our experiments, this integration consistently outperforms either feature type alone, confirming that generative and discriminative representations capture distinct visual information because of their fundamentally different origin.

3.4 Task-Specific Decoders

Our framework uses lightweight task-specific decoders that map extracted features to dense predictions. Decoders are trained on a small set of real annotations before being applied to decode any target modality during generation.

For semantic segmentation, we use the DeepLabV3+ decoder head [6] with ASPP [5] modules. Depth estimation employs a DPT-style decoder [35] with hierarchical reassembly. Salient object detection uses a DPT decoder with a single-channel output. These decoders total approximately 36M trainable parameters, which is much less than a typical DiT (FLUX’s 12B parameters remain frozen). We use frozen FLUX.1-dev [27] as our diffusion backbone and extract features from single-stream blocks at layers $\mathcal{L} = \{4, 15, 26, 37\}$ and timesteps $\mathcal{T} = \{0.78, 0.52, 0.26, 0.001\}$. For decoder training, we use AdamW [29] with a learning rate of 10^{-4} and differential learning rates, where decoders use half the base rate. We apply Exponential Moving Average [34] with $\beta = 0.999$ for stability.

4 Experimental Evaluation

We extensively evaluate MMDiff by training lightweight decoders for three dense prediction tasks: salient object detection, semantic segmentation, and depth estimation. U-Net-based approaches use different generators, meaning both task decoders and generated images differ across methods, making direct comparison of annotations against shared ground truth impossible. Since a single protocol cannot fairly evaluate all aspects simultaneously, we structure our evaluation around three complementary protocols, each addressing a distinct aspect of the pipeline: the consistency of generated image-annotation pairs, the effectiveness of the generated data for training discriminative models, and evaluation on real images against ground-truth labels. Finally, we systematically ablate each component of our framework to quantify its individual contributions to overall performance.

Table 1: Multi-Modal Consistency. Each method generates images with their pipeline and extracts other modalities. We evaluate annotation accuracy against manually verified reference labels on each method’s generated images.

Semantic Segmentation								
Method	mIoU $\uparrow$	Saliency Detection			Depth Estimation			
		Method	$S_m \uparrow$	$F_m \uparrow$	MAE $\downarrow$	AbsRel $\downarrow$	RMSE $\downarrow$	$\delta_1 \uparrow$ $\delta_2 \uparrow$
DatasetDM	62.56	DatasetDM	0.922	0.956	0.038	2.408	0.441	0.151 0.277
DiffuMask	57.35	MMDiff (Ours)	0.944	0.966	0.019	1.359	0.229	0.247 0.469
Dataset Diff.	57.50							
MMDiff (Ours)	72.25							

4.1 Experimental Setup

We evaluate MMDiff across three dense prediction tasks: semantic segmentation, salient object detection, and depth estimation. We report mean Intersection over Union (mIoU) for segmentation, S-measure (S_m) [13], max F-score (F_m) [1]; MAE [33] for saliency; AbsRel, RMSE, and threshold accuracies δ_i for depth.

We compare against diffusion-based feature extractors (VPD [54], Diffusion Hyperfeatures [30], DatasetDM [46]), discriminative encoders (DINO-v3 [38]), and synthetic data generation methods (DiffuMask [47], Dataset Diffusion [31]). To ensure fair comparison across all methods in Table 4, we retrain all baseline methods under identical conditions. Each baseline uses the same training protocol (20 epochs, the same optimizer and learning rate), the same task-specific decoder architectures, and the same data augmentation strategy. This controlled setup isolates the impact of feature quality from differences in training procedures or decoder architectures. The complete pipeline, including training decoders on real data, FLUX synthetic generation with label extraction, and downstream model training on synthetic data, requires approximately 24 hours per task on two 48GB GPUs.

4.2 Multi-Modal Consistency Evaluation

A key challenge in evaluating multi-modal diffusion models is that the generated images differ across methods, making a shared ground truth impossible. We instead measure consistency between each method’s generated images and their predicted modalities, using SOTA models as pseudo-GT. Additionally, we manually verify the pseudo-GT annotations for correctness. For each task, we generate 200 images using the same text prompts across all methods. Each method generates images using its own generation pipeline and extracts annotations using its own approach. We then create reference annotations for each method’s generated images: For salient object detection, we run S3OD [25] on each method’s generated images and manually correct any errors. For depth, we use Depth Anything v2 [49]. For semantic segmentation, we run EoMT [21] and manually correct any errors in the resulting masks.

Table 1 reports multi-modal consistency against pseudo-GT labels comparing MMDiff to Stable Diffusion-based multi-modal generation approaches [31, 46,

Table 2: Semantic segmentation on PASCAL VOC 2012. Our method outperforms DatasetDM [46], DiffuMask [47], and Dataset Diffusion [31] when trained purely on synthetic data, and achieves the best performance when combining synthetic pre-training with real fine-tuning.

Syn.	Real										mIoU
X	✓	91.5	45.7	93.1	88.6	91.9	89.3	91.1	85.0		84.8
DatasetDM [46]	X	87.6	33.8	88.5	74.4	85.7	44.3	83.4	47.7		65.9
DiffuMask [47]	X	84.7	35.8	70.3	52.5	53.1	52.8	75.8	42.9		54.7
Dataset Diffusion [31]	X	79.6	40.2	86.5	69.2	82.7	78.7	69.2	45.3		68.2
DINO-v3 [38]	X	88.9	41.0	90.4	78.5	82.6	86.7	84.3	63.2		75.2
MMDiff (Ours)	X	93.6	45.1	88.8	85.4	83.9	92.3	89.8	69.1		78.9
DatasetDM [46]	✓	95.7	47.9	96.0	92.5	95.2	92.9	92.1	84.2		85.2
DiffuMask [47]	✓	96.4	47.1	97.0	92.7	93.9	93.1	92.9	84.3		86.7
Dataset Diffusion [31]	✓	95.6	47.5	95.7	91.4	94.1	92.6	95.8	84.0		85.7
DINO-v3 [38]	✓	94.2	46.1	96.3	90.7	94.0	92.0	93.7	83.0		85.3
MMDiff (Ours)	✓	96.4	47.7	97.2	89.6	95.7	93.5	95.5	87.3		87.8

47]. Our method achieves substantially higher accuracy than baselines across all metrics. This demonstrates that our DiT-based features more accurately capture the visual content of generated images compared to U-Net-based extraction from Stable Diffusion.

4.3 Synthetic Data Generation

We further validate MMDiff by evaluating the pipeline’s effectiveness as a synthetic data generator. Unlike the consistency evaluation, which evaluates MMDiff using diffusion features only, we use the full model, combining diffusion and DINO-v3 features, to achieve optimal performance. This also enables a direct comparison with DINO-v3 used as a standalone predictor, which serves as a natural baseline: one could simply generate images with FLUX and extract labels using a DINO-based model. We show that all prior methods fail to outperform this baseline, and the combination of discriminative and generative features we extract further improves performance.

We use EoMT [21], a state-of-the-art encoder-only mask transformer, as the downstream model. Pure synthetic training achieves 78.9% mIoU, outperforming all prior synthetic generation methods: DatasetDM (65.9%), DiffuMask (54.7%) and Dataset Diffusion (68.2%) all fall below the simple DINO-v3 standalone baseline (75.2%). This confirms that simply running a discriminative model on generated images outperforms prior UNet-based methods. In contrast, MMDiff shows that diffusion and discriminative features are complementary, outperforming all baselines and demonstrating the approach’s effectiveness for synthetic data generation.

Table 3: Synthetic data evaluation across tasks. We compare DatasetDM, DINO and MMDiff under synthetic-only training, real-only training, and synthetic pre-training followed by real fine-tuning across salient object detection (DUTS) and depth estimation (NYU Depth v2).

Syn.	Real	$S_m \uparrow$	$F_m \uparrow$	MAE $\downarrow$	AbsRel $\downarrow$	RMSE $\downarrow$	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
X	✓	0.838	0.847	0.058	0.1291	0.416	0.868	0.983	0.996
DatasetDM [46]	X	0.781	0.788	0.084	0.3918	1.044	0.620	0.879	0.957
DINO-v3 [38]	X	0.764	0.787	0.109	0.2389	0.626	0.647	0.887	0.972
MMDiff (Ours)	X	0.784	0.798	0.070	0.1880	0.521	0.714	0.914	0.982
DatasetDM [46]	✓	0.813	0.840	0.076	0.1648	0.490	0.804	0.982	0.996
DINO-v3 [38]	✓	0.849	0.853	0.053	0.1415	0.452	0.850	0.987	0.997
MMDiff (Ours)	✓	0.863	0.876	0.049	0.1185	0.392	0.885	0.989	0.998

Table 4: Feature Quality Validation: MMDiff combines features from multiple sources, outperforming both existing diffusion-based extractors and discriminative encoders such as DINO-v3 [38] across all tasks. All methods use frozen encoders except where noted.

Method	Trainable PASCAL VOC 2012		DUTS			NYU Depth v2				
	Encoder	mIoU $\uparrow$	$S_m \uparrow$	$F_m \uparrow$	MAE $\downarrow$	AbsRel $\downarrow$	RMSE $\downarrow$	$\delta_1 \uparrow$	$\delta_2 \uparrow$	$\delta_3 \uparrow$
DatasetDM [46]	X	41.19	0.845	0.909	0.077	0.1536	0.612	0.772	0.971	0.983
Diffusion Hyperfeatures [30]	X	67.57	0.892	0.910	<u>0.021</u>	0.1348	0.503	0.844	0.971	0.978
VPD [54]	✓	82.36	0.912	0.916	0.024	0.1244	<u>0.402</u>	0.867	0.973	0.981
MMDiff (Ours)	X	<u>78.90</u>	0.918	0.929	0.020	0.1175	0.370	0.951	0.976	0.983
DINO-v3 [38]	✓	83.09	<u>0.920</u>	<u>0.922</u>	0.021	0.1288	0.389	0.945	<u>0.975</u>	0.983
VPD + DINO-v3 [38, 54]	✓	<u>83.77</u>	0.918	0.914	<u>0.020</u>	<u>0.1223</u>	<u>0.387</u>	<u>0.946</u>	<u>0.975</u>	0.983
MMDiff + DINO-v3	X	84.95	0.934	0.947	0.018	0.1164	0.365	0.952	0.976	0.983

Table 3 reports results for salient object detection on DUTS and depth estimation on NYU Depth v2. Consistent with segmentation, all prior methods fall below the DINO-v3 standalone baseline in the synthetic-only setting, while MMDiff outperforms it across both tasks. Combining synthetic pretraining with fine-tuning on real data yields the best results across all metrics, confirming that labels extracted by our pipeline complement real data.

4.4 Evaluation on Real Images

Finally, we fix the image and evaluate the predicted modalities against the ground-truth labels, isolating prediction quality from generation. For real images, we follow the noising procedure described in Section 3 to extract features from the frozen backbone. We generate conditioning prompts by captioning each image with Gemma-3-4B [15]. This evaluation addresses two questions: how do frozen DiT features compare with discriminative encoders and methods that fine-tune diffusion backbones for discriminative tasks, and what is the impact of multi-timestep aggregation versus optional DINO integration? Table 4 evaluates the quality of extracted features by comparing performance across three




Fig. 4: Qualitative comparisons on PASCAL (top), DUTS (middle), and NYU (bottom). Our method achieves sharper boundaries and more consistent depth ordering than prior works.

dense prediction tasks. This controlled evaluation demonstrates that our multi-timestep aggregation, concept-driven attention, and discriminative feature integration produce strong representations for dense prediction.

Comparison with U-Net-based Diffusion Methods. MMDiff substantially outperforms all U-Net-based diffusion methods while keeping the encoder fully frozen. Compared with VPD [54], which requires fine-tuning the denoising UNet, we achieve a +2.08% mIoU on segmentation and a +0.011 S-measure on saliency with a frozen encoder. Diffusion Hyperfeatures [30], another U-Net extractor, achieves only 67.57% mIoU compared to our 84.44%. The comparison with Diffusion Hyperfeatures is particularly informative, as both methods extract multi-timestep features, yet we achieve a 16.87% higher mIoU. This validates that our adaptive aggregation mechanism, which utilizes spatially varying weights, is crucial for effectively combining temporal information from transformer architectures. Notably, our FLUX-only features (78.90% mIoU) outperform Diffusion Hyperfeatures (67.57%) by 11.33%, demonstrating the superiority of transformer-based diffusion architectures over U-Nets for feature extraction.

Frozen vs. Trainable Encoders. An important distinction is that our method achieves 84.44% mIoU with a completely frozen 12B-parameter FLUX encoder, whereas VPD (82.36%) and DINO-v3 (83.09%) require encoder fine-tuning. This demonstrates that our features are sufficient to support dense prediction without requiring expensive encoder adaptation. For synthetic data generation, this frozen approach is essential: we can generate annotations using the same frozen features across all tasks without task-specific fine-tuning.

Feature Complementarity. While our frozen diffusion features alone are competitive with finetuned discriminative encoders, integrating DINO-v3 further improves performance from 78.90% to 84.44% mIoU. Crucially, this combination

Table 5: Ablation study. Single-stream features outperform double-stream features even at single timestep. Multi-timestep aggregation provides the largest improvement. Concept-driven attention and DINO integration provide complementary benefits, with the full model achieving best performance across all tasks.

Component	VOC'12	DUTS			NYU Depth v2				
	mIoU $\uparrow$	S_m $\uparrow$	F_m $\uparrow$	MAE $\downarrow$	AbsRel $\downarrow$	RMSE $\downarrow$	δ_1 $\uparrow$	δ_2 $\uparrow$	δ_3 $\uparrow$
double-stream block ($t=0$)	40.52	0.884	0.916	0.033	0.622	1.213	0.740	0.930	0.975
single-stream block ($t=0$)	50.20	0.890	0.914	0.032	0.478	0.738	0.780	0.945	0.978
+ Cross-attention maps	58.31	0.869	0.900	0.043	0.455	0.723	0.800	0.953	0.980
+ Concept attention	63.77	0.903	0.921	0.024	0.431	0.700	0.820	0.960	0.982
Multi-timestep + Concept attention	78.90	0.918	0.929	0.020	0.118	0.370	0.951	0.976	0.983
Full model (w/o aggregation)	78.21	0.918	0.927	0.017	0.130	0.571	0.814	0.972	0.979
Full model (4T + Concepts + DINO)	84.95	0.934	0.947	0.018	0.116	0.365	0.952	0.976	0.983

outperforms both FLUX-only (78.90%) and DINO-only (83.09%) features, as well as VPD+DINO-v3 (83.77%), thereby confirming that our multi-timestep diffusion features capture complementary information to that of discriminative encoders. Diffusion features encode rich semantic and compositional knowledge from their generative training objective, while DINO captures distinct visual characteristics from contrastive learning. Together, these results establish three findings: for the first time, we show that frozen DiT features are competitive with fine-tuned discriminative encoders; multi-timestep aggregation is the primary driver of this performance; and combining generative and discriminative features yields the best results because of their fundamentally different training objectives.

Qualitative Results. Figure 4 demonstrates the superior quality of our extracted features across semantic segmentation, salient object detection, and depth estimation. Our method produces coherent multi-object segmentations, clean saliency maps with precise boundaries, and smooth depth predictions with preserved structural detail. In contrast, U-Net-based methods exhibit semantic confusion, fragmented predictions, and noisy outputs.

4.5 Ablation Study

Table 5 evaluates the contribution of each component in our framework. We systematically ablate components to quantify their impact on the pipeline.

Multi-Timestep Token Aggregation: Aggregating features over four timesteps yields the largest improvement, increasing mIoU from 50.20% to 78.90% (+28.7%). This substantial gain validates our core hypothesis: semantic information in diffusion transformers is temporally distributed across the denoising trajectory. Early timesteps capture global structure and scene layouts, while later timesteps refine local details. Single-timestep extraction at $t = 0$ captures only the final refined features, missing the coarse semantic information present at earlier steps.

Concept-Driven Attention: Concept attention provides substantial improvements with single-timestep features (+13.57% mIoU), demonstrating its value

for spatial guidance. While temporal aggregation captures much of the semantic structure across the denoising trajectory, concept maps contribute task-specific spatial cues that enhance prediction quality. Concept attention also improves upon averaged cross-attention maps, achieving 63.77% mIoU compared to 58.31% with single-timestep features. This demonstrates that querying specific semantic concepts produces sharper, more task-relevant spatial guidance than averaging across all text tokens.

Discriminative Feature Integration: Adding DINO-v3 features to our multi-timestep FLUX features improves the mIoU from 78.90% to 84.44% (+5.54%). This demonstrates that combining features from generative and discriminative pre-training objectives yields complementary information, with the final model outperforming both FLUX-only (78.90%) and DINO-only (83.09%) models.

These ablations explain why our synthetic data outperforms prior methods. Multi-timestep aggregation captures the full semantic trajectory, producing more consistent annotations across diverse scenes. Concept-driven attention provides task-specific spatial guidance, yielding sharper boundaries in the extracted masks. Together, these components enable high-quality label extraction, making our synthetic data practical for training discriminative models.

5 Extension to Video Diffusion Models

To test whether feature aggregation transfers beyond image-space dense prediction, we apply MMDiff to a video-to-3D setting. We adopt the latent-to-3DGS [20] pipeline of Lyra [3] and Wonderland [28], which reconstructs 3D Gaussians from a denoised video latent of world-consistent video model [2, 10, 26, 52], and replace that input with our multi-timestep aggregated features. We train the Lyra decoder on 1,000 RealEstate10K [56] videos and evaluate on 200 held-out scenes.

Table 6: Video $\rightarrow$ 3DGS transfer. Replacing the denoised latent with aggregated MMDiff features yields consistent gains in a low-data setting.

Input	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Denoised Latent Only	21.37	0.705	0.258
MMDiff Aggregated Features	21.61	0.747	0.241

Aggregated features improve over the denoised latent across all three metrics (+0.24 PSNR, +0.042 SSIM, -0.017 LPIPS), indicating that the benefit of multi-timestep aggregation is not specific to image-space dense prediction but extends to video and 3D. Fig. 5 shows a layer/timestep analysis on Wan-2.1 [43], which follows the same trends observed for image DiTs: early timesteps and shallow layers carry fine detail, later timesteps and deeper layers carry semantic structure.




Fig. 5: Wan-2.1 layer/timestep analysis. Feature behavior across video DiT layers and timesteps mirrors the image-model trends.

6 Conclusion

We presented MMDiff, a framework that transforms a frozen diffusion transformer into a multi-modal generative system that extracts dense perceptual annotations during image generation. By recycling the rich features computed across the denoising trajectory rather than discarding them after rendering, our approach addresses a key inefficiency in current generative pipelines: the need to re-extract semantic, geometric, and structural information from generated pixels using separate task-specific models. Our work demonstrates that transformer-based diffusion models encode substantially richer perceptual representations than U-Net architectures. This architectural advantage stems from the temporal distribution of semantic information across the denoising trajectory, which we capture through adaptive multi-timestep aggregation with spatially varying weights. Combined with concept-driven attention extraction for targeted spatial guidance and optional discriminative feature integration, this approach yields high-quality annotations that substantially outperform prior U-Net-based methods on synthetic data generation, direct annotation quality, and real-image prediction tasks. By enabling multi-modal annotation extraction from a single frozen generative model, MMDiff provides a practical framework for cost-effective, scalable synthetic dataset creation. This opens new possibilities for leveraging generative models in vision tasks where manual annotation is expensive or prohibitive, and establishes that diffusion transformers can serve as effective feature extractors beyond their generative capabilities.

Acknowledgments

This research was funded by the European Union (ERC, Volute, 101222037). O.K. is supported by a Google unrestricted gift. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

References

1. Achanta, R., Hemami, S., Estrada, F., Susstrunk, S.: Frequency-tuned salient region detection. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 1597–1604. IEEE (2009)
2. An, Z., Kupyn, O., Uscidda, T., Colaco, A., Ahuja, K., Belongie, S., Gonzalez-Franco, M., Gazulla, M.T.: Vggrpo: Towards world-consistent video generation with 4d latent reward. arXiv preprint arXiv:2603.26599 (2026)
3. Bahmani, S., Shen, T., Ren, J., Huang, J., Jiang, Y., Turki, H., Tagliasacchi, A., Lindell, D.B., Gojcic, Z., Fidler, S., Ling, H., Gao, J., Ren, X.: Lyra: Generative 3d scene reconstruction via video diffusion model self-distillation. In: International Conference on Learning Representations (ICLR) (2026)
4. Chen, J., YU, J., GE, C., Yao, L., Xie, E., Wang, Z., Kwok, J., Luo, P., Lu, H., Li, Z.: Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In: The Twelfth International Conference on Learning Representations (2024)
5. Chen, L.C., Papandreou, G., Schroff, F., Adam, H.: Rethinking atrous convolution for semantic image segmentation. arXiv preprint arXiv:1706.05587 (2017)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018)
7. Dai, Y., Jiang, F., Wang, C., Xu, M., Qi, Y.: Fantasyworld: Geometry-consistent world modeling via unified video and 3d prediction. arXiv preprint (2025)
8. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
9. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021)
10. Du, H., Ye, J., Cong, X., Li, R., Ni, J., Agarwal, A., Zhou, Z., Li, Z., Balestrieri, R., Wang, Y.: Videogpa: Distilling geometry priors for 3d-consistent video generation. arXiv preprint arXiv:2601.23286 (2026)
11. Erel, Y., Dinkel, O., Dabral, R., Golyanik, V., Theobalt, C., Bermano, A.H.: Attention (as discrete-time markov) chains. arXiv preprint arXiv:2507.17657 (2025)
12. Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Rombach, R.: Scaling rectified flow transformers for high-resolution image synthesis. In: Forty-first International Conference on Machine Learning (2024)
13. Fan, D.P., Cheng, M.M., Liu, Y., Li, T., Borji, A.: Structure-measure: A new way to evaluate foreground maps. In: Proceedings of the IEEE international conference on computer vision. pp. 4548–4557 (2017)
14. Gan, C., Tu, Y., Chen, X., Chen, T., Li, Y., Harandi, M., Lin, W.: Unleashing diffusion transformers for visual correspondence by modulating massive activations. arXiv preprint arXiv:2505.18584 (2025)
15. Gemma Team: Gemma 3 (2025), <https://goole/Gemma3Report>, accessed: 2026-06-25
16. Helbling, A., Meral, T.H.S., Hoover, B., Yanardag, P., Chau, D.H.: Conceptattention: Diffusion transformers learn highly interpretable features (2025)
17. Hendrycks, D.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)

18. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
19. Huang, C.H.P., Mitra, N., Jeong, H., Yoon, J.S., Ceylan, D.: Jog3r: Towards 3d-consistent video generators. *arXiv preprint* (2025)
20. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics* **42**(4) (July 2023)
21. Kerssies, T., Cavagnero, N., Hermans, A., Norouzi, N., Averta, G., Leibe, B., Dubbelman, G., de Geus, D.: Your vit is secretly an image segmentation model. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 25303–25313 (2025)
22. Kim, C., Shin, H., Hong, E., Yoon, H., Arnab, A., Seo, P.H., Hong, S., Kim, S.: Seg4diff: Unveiling open-vocabulary segmentation in text-to-image diffusion transformers. *arXiv preprint arXiv:2509.18096* (2025)
23. Krishnan, A., Yan, X., Casser, V., Kundu, A.: Orchid: Image latent diffusion for joint appearance and geometry generation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 28217–28227 (2025)
24. Kukleva, A., Simsar, E., Tonioni, A., Naeem, M.F., Tombari, F., Lenssen, J.E., Schiele, B.: Refam: Attention magnets for zero-shot referral segmentation. *arXiv preprint arXiv:2509.22650* (2025)
25. Kupyn, O., Kataoka, H., Rupprecht, C.: S3od: Towards generalizable salient object detection with synthetic data. In: *International Conference on Learning Representations (ICLR)* (2026)
26. Kupyn, O., Manhardt, F., Tombari, F., Rupprecht, C.: Epipolar geometry improves video generation models. *arXiv preprint arXiv:2510.21615* (2025)
27. Labs, B.F.: Flux. <https://github.com/black-forest-labs/flux>, Accessed: 2026-06-25 (2023)
28. Liang, H., Cao, J., Goel, V., Qian, G., Korolev, S., Terzopoulos, D., Plataniotis, K.N., Tulyakov, S., Ren, J.: Wonderland: Navigating 3d scenes from a single image. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 798–810 (2025)
29. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
30. Luo, G., Dunlap, L., Park, D.H., Holynski, A., Darrell, T.: Diffusion hyperfeatures: Searching through time and space for semantic correspondence. *Advances in Neural Information Processing Systems* **36**, 47500–47510 (2023)
31. Nguyen, Q., Vu, T., Tran, A., Nguyen, K.: Dataset diffusion: Diffusion-based synthetic data generation for pixel-level semantic segmentation. *Advances in Neural Information Processing Systems* **36**, 76872–76892 (2023)
32. Peebles, W., Xie, S.: Scalable diffusion models with transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 4195–4205 (2023)
33. Perazzi, F., Krähenbühl, P., Pritch, Y., Hornung, A.: Saliency filters: Contrast based filtering for salient region detection. In: *2012 IEEE conference on computer vision and pattern recognition*. pp. 733–740. IEEE (2012)
34. Polyak, B.T., Juditsky, A.B.: Acceleration of stochastic approximation by averaging. *SIAM journal on control and optimization* **30**(4), 838–855 (1992)
35. Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 12179–12188 (2021)

36. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
37. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. pp. 234–241. Springer (2015)
38. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., et al.: Dinov3. *arXiv preprint arXiv:2508.10104* (2025)
39. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations* (2021)
40. Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R.: Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research* **15**(1), 1929–1958 (2014)
41. Tang, L., Jia, M., Wang, Q., Phoo, C.P., Hariharan, B.: Emergent correspondence from image diffusion. *Advances in Neural Information Processing Systems* **36**, 1363–1389 (2023)
42. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 32–42 (2021)
43. Wan, T., Wang, A., Ai, B., Wen, B., Mao, C., Xie, C.W., Chen, D., Yu, F., Zhao, H., Yang, J., Zeng, J., Wang, J., Zhang, J., Zhou, J., Wang, J., Chen, J., Zhu, K., Zhao, K., Yan, K., Huang, L., Feng, M., Zhang, N., Li, P., Wu, P., Chu, R., Feng, R., Zhang, S., Sun, S., Fang, T., Wang, T., Gui, T., Weng, T., Shen, T., Lin, W., Wang, W., Wang, W., Zhou, W., Wang, W., Shen, W., Yu, W., Shi, X., Huang, X., Xu, X., Kou, Y., Lv, Y., Li, Y., Liu, Y., Wang, Y., Zhang, Y., Huang, Y., Li, Y., Wu, Y., Liu, Y., Pan, Y., Zheng, Y., Hong, Y., Shi, Y., Feng, Y., Jiang, Z., Han, Z., Wu, Z.F., Liu, Z.: Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314* (2025)
44. Wang, S., Leroy, V., Cabon, Y., Chidlovskii, B., Revaud, J.: Dust3r: Geometric 3d vision made easy. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 20697–20709 (2024)
45. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: Cbam: Convolutional block attention module. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 3–19 (2018)
46. Wu, W., Zhao, Y., Chen, H., Gu, Y., Zhao, R., He, Y., Zhou, H., Shou, M.Z., Shen, C.: Datasetdm: Synthesizing data with perception annotations using diffusion models. *Advances in Neural Information Processing Systems* **36**, 54683–54695 (2023)
47. Wu, W., Zhao, Y., Shou, M.Z., Zhou, H., Shen, C.: Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 1206–1217 (2023)
48. Xu, G., Ge, Y., Liu, M., Fan, C., Xie, K., Zhao, Z., Chen, H., Shen, C.: What matters when repurposing diffusion models for general dense perception tasks? In: *The Thirteenth International Conference on Learning Representations* (2025)
49. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. *Advances in Neural Information Processing Systems* **37**, 21875–21911 (2024)

50. Yue, Z., Wang, J., Sun, Q., Ji, L., Chang, E.I., Zhang, H., et al.: Exploring diffusion time-steps for unsupervised representation learning. arXiv preprint arXiv:2401.11430 (2024)
51. Zhang, Q., Zhai, S., Martin, M.A.B., Miao, K., Toshev, A., Susskind, J., Gu, J.: World-consistent video diffusion with explicit 3d modeling. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 21685–21695 (2025)
52. Zhang, Q., Zhai, S., Martin, M.A.B., Miao, K., Toshev, A., Susskind, J., Gu, J.: World-consistent video diffusion with explicit 3d modeling. In: CVPR. pp. 21685–21695 (2025)
53. Zhang, Y., Wang, Z., Zheng, W., Zhu, Z., Zhou, J., Lu, J.: Joint 3d geometry reconstruction and motion generation for 4d synthesis from a single image. arXiv preprint arXiv:2512.05044 (2025)
54. Zhao, W., Rao, Y., Liu, Z., Liu, B., Zhou, J., Lu, J.: Unleashing text-to-image diffusion models for visual perception. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5729–5739 (2023)
55. Zheng, S., Bao, Z., Zhao, R., Hebert, M., Wang, Y.X.: Diff-2-in-1: Bridging generation and dense perception with diffusion models. In: The Thirteenth International Conference on Learning Representations (2025)
56. Zhou, T., Tucker, R., Flynn, J., Fyffe, G., Snavely, N.: Stereo magnification: Learning view synthesis using multiplane images. In: SIGGRAPH (2018)