arXiv:2611.24738v1 [cs.CV] 18 Nov 2026

Jeveration: System One Models Now Generate

Yuhe Liu* Qi Li* Shihua Zhang Xinchao Wang†
National University of Singapore
yuhe_liu@u.nus.edu liqi@u.nus.edu suhzhang001@gmail.com xinchao@nus.edu.sg

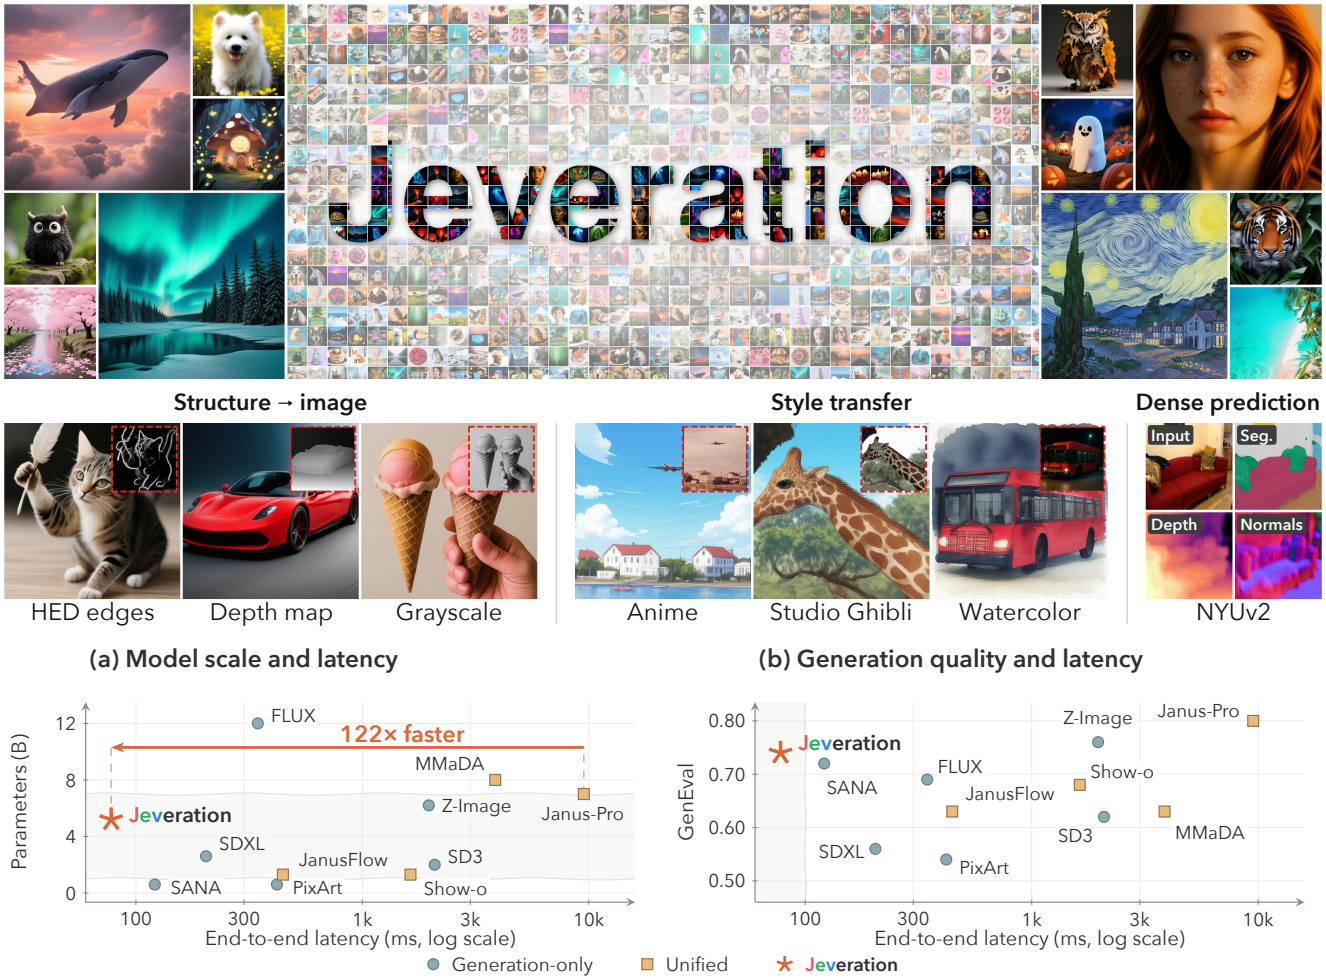


Figure 1. **System One models now generate.** One frozen vision-language model regards text-to-image and image-to-image tasks as typed decisions. Text-to-image samples at 512 px (top) and image-conditioned generation and dense prediction (middle), all from the same Jeveration checkpoint. The bottom charts plot model size (a) and GenEval (b) against latency on one H200.

Abstract

System One models such as Jev return decisions through fast, parallel prediction. In this paper, we introduce Jeveration, showing that the same decision paradigm can also generate high-quality images in a single forward pass. Jeveration builds on a frozen vision-language model, repurposing its visual-semantic knowledge for holistic image generation. The key challenge is to maintain global consistency. A single prompt admits many valid images, yet parallel predictions must converge on a coherent whole. We therefore condition all positions on a shared random state to coordinate parallel predictions, so that each draw yields a different but coherent image. Since the model tends to ignore this state under direct image supervision, we first learn a distribution over discrete image codes to guide single-pass gen-

*Equal contribution †Corresponding author

eration. Beyond text-to-image generation, we further unify image-to-image generation and dense prediction within the same architecture, so that one checkpoint supports all three tasks, each in one forward pass. Experiments show that Jeveration generates a 512×512 *image in 78 ms, up to* $122\times$ *faster than existing unified models, while maintaining competitive image quality and dense prediction performance.*

1. Introduction

Image generation today is dominated by two paradigms. Autoregressive models [11, 40, 44] generate an image token by token, while diffusion models [8, 25, 48] generate it by iterative denoising. Both generally require many sequential model evaluations for a single image, which increases inference latency. We introduce a third paradigm for image generation that removes sequential sampling, built on System One models. Our model, Jeveration, returns a complete image as one joint decision in a single forward pass, given a text prompt or a source image (Fig. 1).

System One models such as Jev [1] return typed decisions, such as Choices or Scores, in a single parallel query without extended reasoning [1]. They have been applied to option selection, object localization, robot action scoring, and image evaluation [12, 26, 27, 42, 43]. Their effectiveness in these decision-making and understanding tasks draws in part on the visual and semantic knowledge acquired through large-scale pretraining. This knowledge also offers a foundation for image synthesis, suggesting that System One models have generative potential beyond their current uses. Motivated by this potential, we first explore a direct approach: using a System One model to select an entry from an existing codebook for every code position in a single forward pass. Our initial attempts, however, reveal two key challenges. First, direct parallel selection does not ensure global coherence. A single prompt admits many valid images, thus parallel decisions at different positions may correspond to incompatible realizations, producing image regions that fail to fit together, as illustrated in Fig. 2. Second, extending this approach to diverse image generation tasks requires a unified architecture that accommodates different inputs and conditioning requirements.

To keep the image globally consistent, we condition all positions on a shared random state that coordinates the parallel predictions, so that each draw yields a different but coherent image. Furthermore, since a model trained on target images alone tends to ignore this state, we first train a model that learns the distribution of images as discrete codes, and then use it to judge the single-pass decisions and guide their training. The same paradigm extends to further tasks whose output is an image. Image-to-image generation uses a source image or visual conditions, such as edges, depth or grayscale, to guide the content and appearance



Figure 2. **Direct parallel selection yields incoherent images.** Codes sampled independently from the first-phase model do not fit together, while Jeveration produces coherent images.

of the generated image. Dense prediction can be regarded as image-to-image generation in a broad sense, where the model generates a map aligned with the input image that specifies a class, depth or surface normal at each pixel. One model can thus cover generation as well as perception.

Putting these together, Jeveration starts from a vision-language model (VLM), Qwen3.5-4B [41], and lets it answer visual tasks with Jev-style decisions. The VLM stays frozen, and Jeveration adds only LoRA adapters [21] and output heads. Each answer, whether an image or a dense map, is one typed decision made in a single forward pass, and one checkpoint covers text-to-image generation, image-to-image generation and dense prediction tasks. It reaches 0.74 GenEval at 512 px in 78 ms per image on one H200, $122\times$ faster than Janus-Pro-7B [11], closely follows the input structure in image-to-image generation, and achieves competitive results on dense prediction. These results show that a decision model itself can generate images and predict dense maps in a single forward pass. Our contributions are:

- **Image generation through typed decisions.** To our knowledge, this is the first work to show that Jev-style system one models can themselves generate images in one forward pass, for both text-to-image and image-to-image generation.
- **Learning coherent image decisions.** We let parallel decisions form one coherent image by conditioning them on one shared random state, and train the model to use this state by first learning the image distribution and then using it to judge and guide the decisions.
- **One checkpoint across visual outputs.** One checkpoint with shared adapters performs text-to-image generation, image-to-image generation and dense prediction. Extensive experiments and ablation studies on these tasks demonstrate its effectiveness and efficiency.

2. Related work

Unified multimodal understanding and generation. Unified multimodal models combine image understanding and synthesis. Chameleon [40], Emu3 [44], Janus/Janus-Pro [11, 45] and Unified-IO 2 [33] generate image tokens autoregressively. Show-o [49] uses masked prediction, while Transfusion [56] and BAGEL [14] model continuous image latents. MetaQueries [36] and BLIP3-o [9] instead

connect understanding representations to separate diffusion generators. All of these models generate an image over many sequential steps, token by token or by iterative denoising. Dense prediction is another route to visual outputs. For instance, LENS [32] adds segmentation to a frozen multimodal model, while DenseMLLM [29] supervises visual tokens for dense prediction. In contrast, Jeveration unifies image synthesis and dense prediction by formulating both as parallel typed decisions, using a single frozen VLM with shared adapters and task-specific output heads.

One-step image generation. Single-pass generators learn to turn a random input into an image without iterative sampling. Consistency models [39] and MeanFlow [18] learn direct generation in continuous spaces. DMD/DMD2 [52, 53], SANA-Sprint [10] and TwinFlow [13] provide further approaches to one-step or few-step synthesis. Acceleration also applies to unified models. For example, Show-o Turbo [50] uses consistency distillation, and Omni-Diffusion-Distill [23] compresses multimodal decoding trajectories. For discrete images, MaskGIT [5], Muse [6] and Meissonic [3] learn masked-code conditionals and sample over multiple rounds. Di[M]O [57] distills such a model into a one-step generator by matching its token-level conditionals with the help of an auxiliary model. Soft-Di[M]O [58] further uses soft embeddings to enable differentiable refinement. Building on this line of work, Jeveration adapts token-level conditional matching to learn coherent image synthesis through the parallel typed decisions of a frozen VLM. Its target conditionals are provided by a masked model first trained on the same frozen backbone, while a confidence-guided correction target further guides parallel code selection.

Typed decision models. Jev [1] returns typed outputs such as Choices, Scores, and probabilities in parallel without generating lengthy text, while its Box interface predicts the four edges of an object's bounding box [42]. Open decision models [2, 12, 19] and a clinical decision model [55] similarly provide structured decision outputs. Jeva [28] explicitly addresses candidate-set scalability, demonstrating full-set decision-making with up to 2,048 candidates in its evaluation. These decision interfaces also support specialized applications: Qwen-Audio-Jev [17] answers typed questions about audio, JevVLA [43] scores and corrects robot actions, and Visual Jev Rewards [26] evaluates generated images to reward a separate generator. Jeveration extends this paradigm to spatially coordinated image generation, predicting all image tokens in parallel within a single forward pass conditioned on a shared random state. The same model also predicts dense maps in one pass, using a class Choice at each position for segmentation and Scores for depth and surface normal estimation.

3. Method

Jeveration enables a VLM to generate images through parallel typed decisions in a single forward pass, while supporting dense prediction within the same framework. As shown in Fig. 3, for image synthesis, these decisions select a codebook entry at each code position, while for dense prediction they produce task-specific classes or scores. A shared random state gives all code positions common stochastic conditioning, so that each draw yields a different but globally coherent image. We first formalize these visual decisions and the role of shared randomness (Sec. 3.1). We then describe how a frozen VLM with shared adapters and task-specific output heads supports these tasks within one checkpoint (Sec. 3.2). Finally, we present the training procedure that teaches the model to produce coherent images through parallel decisions (Sec. 3.3).

3.1. Images as Decisions

Given a condition y, which holds a prompt or an input image, Jeveration answers with typed decisions [1]. An image is a joint *Choice* of discrete codes $c = (c_1, \ldots, c_N)$ of an image tokenizer, segmentation is a *Choice* over semantic classes at each spatial position, and depth and surface normals are real-valued *Scores* at every position. A dense map has one correct answer per input and is decided deterministically, whereas one prompt admits many valid images, so an image decision should be stochastic.

Why a joint decision is hard. Given the conditioning input y, such as a text prompt or a source image with instructions, a direct approach predicts a categorical distribution at each code position and samples the codes independently in parallel:

$$q_{\mathrm{ind}}(c \mid y) = \prod_{i=1}^{N} q_i(c_i \mid y). \tag{1}$$

This factorisation cannot capture dependencies between codes. For a fixed condition y and target distribution $p(c \mid y)$, its best forward-KL approximation satisfies

$$\min_{q_{\mathrm{ind}}} \mathrm{KL}(p \,\|\, q_{\mathrm{ind}}) = \sum_{i=1}^{N} H(c_i \mid y) - H(c \mid y), \tag{2}$$

where H denotes entropy under p given y. The gap is the target's total correlation, which measures how strongly the codes of an image depend on each other. A coherent image needs this dependence, but an independent model matches only the marginals $p(c_i \mid y)$, so its codes may follow different images and fail to fit together (Fig. 2).

A shared random state. Jeveration draws one random state sequence ξ per image and conditions every code po-

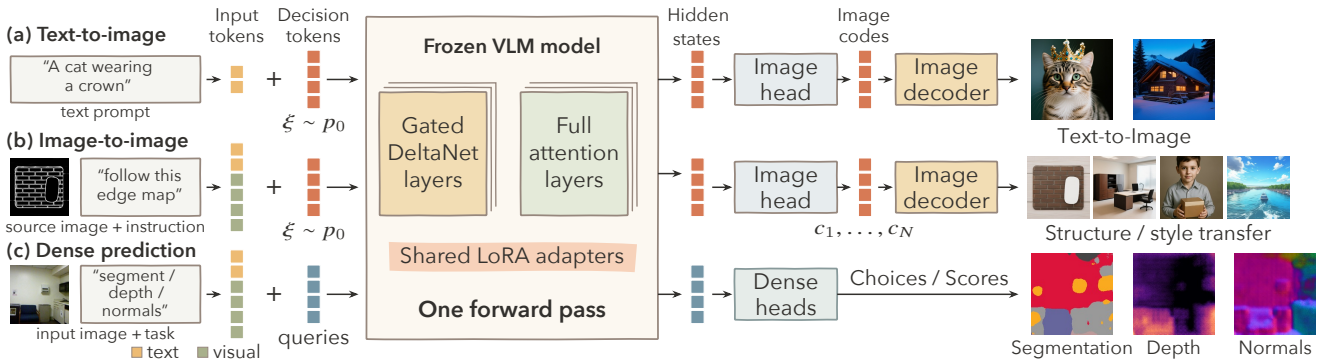


Figure 3. **Architecture overview.** A frozen vision-language model with shared LoRA adapters processes input and decision tokens in one forward pass. Decision tokens start from a random state for image generation, while for dense prediction they act as spatial queries. Task-specific heads then map their final hidden states to image codes or dense predictions. One checkpoint serves all tasks.

sition on this shared state:

$$q_\theta(c \mid y) = \mathbb{E}_{\xi \sim p_0}\left[\prod_{i=1}^{N} q_{\theta,i}(c_i \mid y, \xi)\right]. \tag{3}$$

The shared state ξ provides a common source of randomness for all image positions. Although the codes are sampled independently given (y, ξ), their shared dependence on ξ allows the resulting image distribution to capture dependencies between positions. The model can therefore coordinate its code choices while deciding all positions in a single forward pass. What remains is to learn how to use ξ to produce coherent images (Sec. 3.3).

3.2. Architecture and Inference

Overall architecture. We implement these visual decisions by appending decision tokens to the input of a VLM, as shown in Fig. 3. Text tokens encode the prompt and task instruction, and a frozen vision tower encodes any input image into visual tokens. The decision tokens are initialised from a sampled random state for image generation and serve as spatial queries for dense prediction. A frozen Qwen3.5-4B [41] backbone with shared LoRA adapters [21] processes the input and decision tokens in a single forward pass. In its full-attention layers, decision tokens attend to each other bidirectionally while the input prefix remains causal. Finally, task-specific heads map the decision tokens' final hidden states to image codes or dense predictions. The same checkpoint is used for every task.

Image generation. The image tokenizer is a frozen TA-TiTok [24, 54] with $N = 128$ codes per image. For image generation, we append N decision tokens, and the i-th token decides the code c_i. The initial embeddings of these tokens encode the random state $\xi \sim p_0$, which masks each position with probability ρ, fills the remaining positions with uniformly random codes and adds Gaussian noise of scale σ to all embeddings. Following Di[M]O [57], we set $\rho = 0.6$ and $\sigma = 0.1$. A transformer head [37] reads the final hidden states of these tokens together with ξ and outputs a feature u_i for each position. We predict each code from u_i with two branches. The codebook branch f_{cb} projects u_i to the code space and applies a classifier initialised from the tokenizer's codebook. The wide branch f_{wide}, initialised at zero, maps u_i directly to the code logits, giving the head more freedom to learn. The code distribution sums the two:

$$q_{\theta,i}(\cdot \mid y, \xi) = \text{softmax}\left(f_{\text{cb}}(u_i) + f_{\text{wide}}(u_i)\right). \tag{4}$$

Both branches are trained and we conduct ablations in Sec. 4.4 to show the benefit of combining them. At inference, all codes are sampled in parallel as one joint decision. Finally, they are rendered by the tokenizer's decoder, which also reads the prompt and which we fine-tune for 512-px output. For image-to-image generation, a gated branch also feeds the source image to the decoder.

Dense prediction. For dense prediction, the decision tokens are spatial queries aligned with the visual tokens of the input image. Each query starts from a learned vector plus a projection of its visual token. Task-specific heads decode each query's final hidden state into the segmentation, depth or surface-normal predictions for its image region.

3.3. Learning to Decide

We train the model to use ξ to make coherent image decisions. Direct supervision on target codes does not encourage the model to use ξ, since the random state is sampled independently of the target image. We therefore first learn a model of the image distribution through masked prediction, then use it to guide image decisions from random states.

Learning the image distribution. In the first phase, we train the architecture as a masked model [5, 6], keeping the

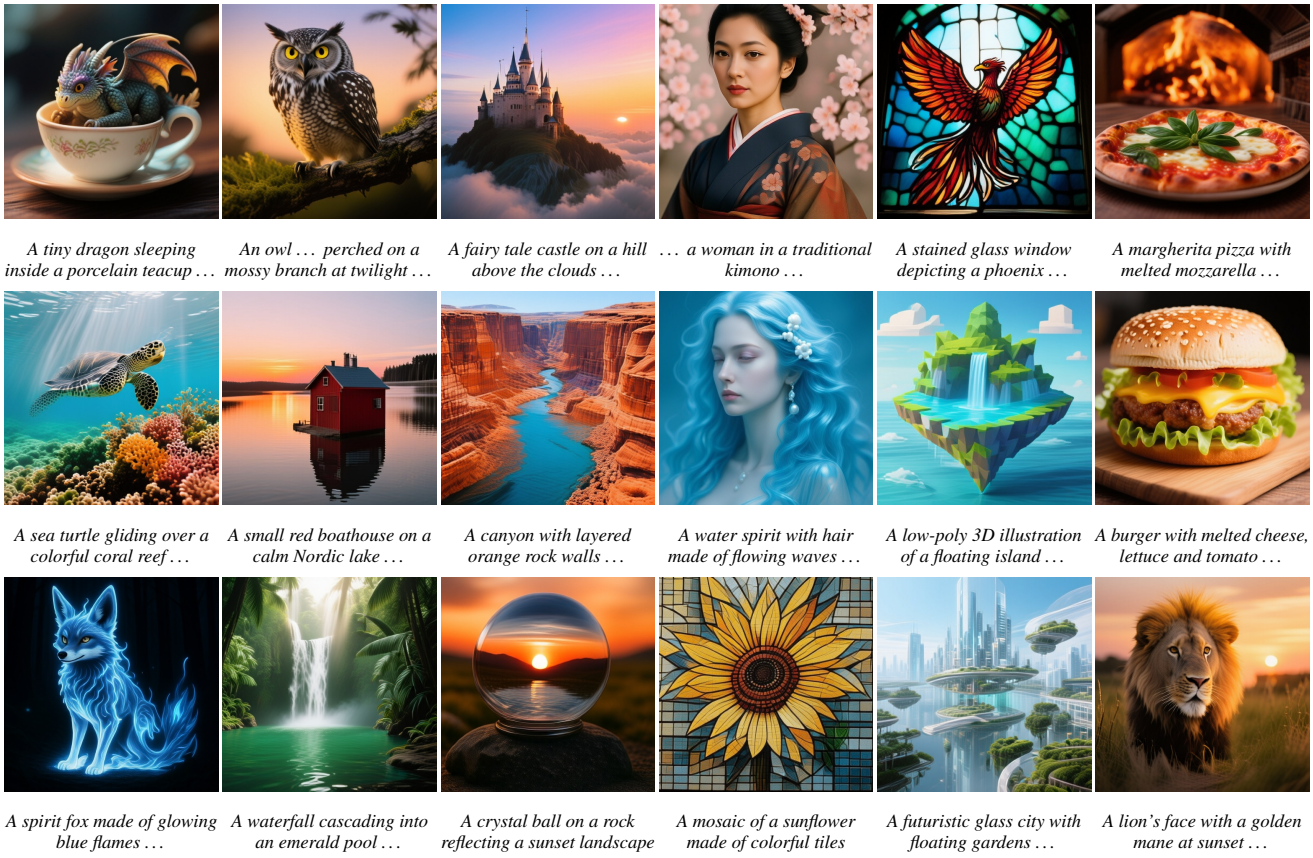


Figure 4. **Text-to-image generation.** Samples of the final Jeveration checkpoint at 512 px, with excerpts of their prompts.

VLM frozen. For target codes c, we mask a random set M and predict the hidden codes from the visible codes $c_{\bar{M}}$ and condition y:

$$\mathcal{L}_{\text{mask}}(\phi) = \mathbb{E}\left[\frac{1}{|M|}\sum_{i\in M} -\log p_{\phi,i}(c_i \mid c_{\bar{M}}, y)\right]. \quad (5)$$

The visible codes supply image-specific context, enabling the model to learn dependencies between positions. The resulting model provides reference conditional predictions for the second phase.

Learning image decisions. In the second phase, the decision model θ starts from the first-phase model and learns to turn random states into whole images. From (y, ξ), θ predicts logits $\ell_{\theta,i}$ and samples a code sequence $\hat{c}$ in a single pass. We re-mask a random subset M of $\hat{c}$, and two masked models predict the hidden codes from the same visible codes. The frozen first-phase model ϕ represents the image distribution and gives $p_{\phi,i}$. An auxiliary model ψ gives $p_{\psi,i}$ for the current distribution of θ. The auxiliary model begins as a copy of ϕ and keeps training on θ's samples with the masked loss of Eq. (5), so it follows θ as θ changes. The two predictions agree when θ matches the image distribution. Following Di[M]O [57], we use their difference to update θ by minimising

$$\mathcal{L}_{\text{DM}}(\theta) = \mathbb{E}\left[\frac{1}{|M|}\sum_{i\in M} \text{sg}(p_{\psi,i} - p_{\phi,i})^\top \ell_{\theta,i}\right], \quad (6)$$

where sg stops gradients. The prediction at each masked position depends on the other codes of the sample, so the two predictions agree only when θ's codes fit together. This pushes θ to coordinate its codes through the shared state ξ.

Confidence-guided correction. In the final part of the second phase, we add direct supervision on generated samples. For some samples in each batch, we re-mask the half of the positions with the lowest confidence $q_{\theta,i}(\hat{c}_i \mid y, \xi)$, denoted R, and let the frozen model complete them over several rounds into c^{ref}. We then train θ to predict these corrected codes from the same (y, ξ):

$$\mathcal{L}_{\text{gen}} = \mathcal{L}_{\text{DM}} + \lambda\,\mathcal{L}_{\text{corr}},$$
$$\mathcal{L}_{\text{corr}} = \mathbb{E}_{\text{selected}}\left[\frac{1}{|R|}\sum_{i\in R} -\log q_{\theta,i}(c_i^{\text{ref}} \mid y, \xi)\right]. \quad (7)$$

Table 1. **Text-to-image quality and end-to-end latency.** All models are timed on the same H200 at batch 1, including decoding, with released samplers at the listed steps and resolutions. Speedup is relative to Janus-Pro-7B, the slowest model listed. Params include our frozen backbone; Steps count sequential generator evaluations, or tokens for autoregressive models. Colored squares identify image-generation families, including distilled variants. † / ‡ indicate scores from the authors / another paper.

Model	Unified	Params (B)	Steps	Res. (px)	Speedup↑	GenEval↑	DPG↑	HPSv2.1↑	Latency (ms)↓
Generation-only models									
■ PixArt-α [8]	×	0.6	20	512	25.1×	0.49	71.6‡	29.98	378.8
■ PixArt-Σ [7]	×	0.6	20	512	22.7×	0.54	79.5‡	30.49	419.5
■ SANA-0.6B [48]	×	0.6	20	512	16.2×	0.65	84.3†	29.86	587.3
■ SANA-Sprint 0.6B [10]	×	0.6	1	1024	78.5×	0.72	78.63	28.74	121.2
■ SDXL-DMD2 [52]	×	2.6	1	1024	46.6×	0.56	74.09	28.20	204.0
■ SDXL-LCM [34]	×	2.6	4	1024	28.2×	0.53‡	71.08	26.99	336.7
■ FLUX.1-schnell [25]	×	12.0	1	1024	27.5×	0.69	85.97	29.63	345.6
■ Z-Image-Turbo [4]	×	6.2	8	1024	4.8×	0.76	84.99	29.76	1,965.2
■ SD3-medium [15]	×	2.0	28	1024	4.6×	0.62‡	84.1‡	29.88	2,086.0
Unified models									
■ JanusFlow-1.3B [35]	✓	1.3	30	384	21.3×	0.63†	80.09†	22.90	445.8
■ Janus-Pro-1B [11]	✓	1.5	576	384	1.3×	0.73	82.63†	27.73	7,593.2
■ Janus-Pro-7B [11]	✓	7.0	576	384	1.0×	0.80	83.91	28.47	9,508.4
■ Show-o Turbo-512 [50]	✓	1.3	4	512	101.0×	0.63	83.35	27.35	94.1
■ Show-o-512 [49]	✓	1.3	50	512	5.8×	0.68	67.27‡	27.85	1,631.0
■ MMaDA-8B [51]	✓	8.0	15	512	2.5×	0.63†	69.97‡	23.74	3,860.9
Ours									
■ **Jeveration**	✓	5.2	1	512	**122.2**×	0.74	78.24	28.06	77.8

■ Continuous diffusion/flow family ■ Masked-token generation ■ Autoregressive generation ■ Parallel typed decisions

The frozen model, the auxiliary model and the correction are used only in training, so inference remains a single forward pass of θ.

Dense supervision. Dense tasks are supervised with ground truth, using cross-entropy for segmentation, an ℓ_1 depth error and $1 - \cos$ for surface normals.

4. Experiments

4.1. Experimental Setup

Data and training. All main results use the same Jeveration checkpoint. Generation training combines FLUX-Reason-6M images [16] and Qwen-Image [46] renderings of text-to-image-2M prompts [59]. Image-to-image training uses derived conditions and synthetic style pairs, while dense tasks use NYUv2 [38]. We train rank-256 LoRA adapters for about 92k steps in the first phase and 6.7k in the second, which take about four days on a single machine with four NVIDIA RTX PRO 6000 GPUs.

Quality evaluation. We evaluate text-to-image generation with GenEval [20], DPG-Bench [22] and HPSv2.1 [47]. Image-to-image outputs are scored by how well they follow the input structure, using edge F1 or depth accuracy. Stylized outputs are compared with paired targets by PSNR. Dense prediction uses the standard NYUv2 metrics. Latency is the median of five end-to-end calls on one H200 at batch 1. Additional details are provided in the supplementary material.

4.2. Text-to-Image Generation

Jeveration produces visually coherent images with detailed textures across natural scenes, portraits, objects and artistic styles (Figs. 1 and 4). The samples capture both the requested content and appearance, from the dragon inside a teacup to the stained-glass phoenix. In one forward pass, it generates a 512×512 image in 77.8 ms end to end on one H200, the fastest in Tab. 1, and reaches 0.74 GenEval.

Among unified models, Jeveration is the fastest and achieves higher GenEval and HPSv2.1 than JanusFlow, Show-o, Show-o Turbo, MMaDA and Janus-Pro-1B. At the same 512-px resolution, it improves GenEval by 0.11 over four-step Show-o Turbo-512 while reducing latency from 94.1 to 77.8 ms. Compared with the autoregressive Janus-Pro-1B and Janus-Pro-7B, it is 98× and 122× faster.

Compared with generation-only models, Jeveration achieves higher GenEval than the multi-step models at the same 512-px resolution (0.74 versus 0.49 to 0.65) while being 5 to 8 times faster. It also achieves higher GenEval than

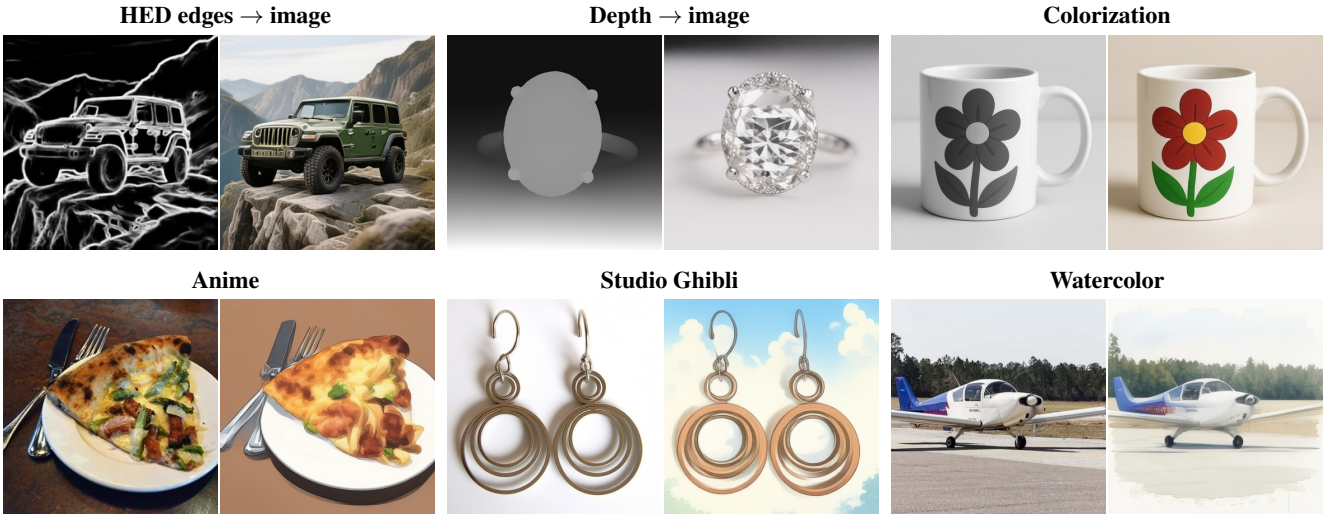


Figure 5. **Image-to-image generation** on held-out examples. In each pair the left image is the input and the right image the output.

Table 2. **Image-to-image generation.** Structure fidelity on 341 held-out examples per task and style PSNR against paired targets.

	Structure follow↑			Style
Model	**HED**	**Depth**	**Colorize**	**PSNR↑**
Jeveration	**0.880**	**0.802**	0.937	**15.40**
Copy the input	0.828	0.713	**0.998**	–
Decoder without source	0.792	0.767	0.757	14.37

Table 3. **Dense prediction on NYUv2.** Dense-only models use the same backbone. ‡ is a reported result.

	Seg.	Depth		Normal
Model	**mIoU↑**	**abs. (m)↓**	**δ_1 ↑**	**mean (°)↓**
Jeveration	67.7	0.267	0.898	24.2
Dense-only, one task each	**67.8**	**0.246**	**0.915**	**21.2**
Dense-only, three tasks	67.3	0.253	0.912	23.0
DeepLabV3+ R50‡ [30, 31]	53.5	0.393	–	22.0

SANA-Sprint 0.6B, FLUX.1-schnell and SDXL-DMD2, one-step models that generate 1024-px images. These results show that a frozen VLM can generate images of competitive quality in one forward pass.

4.3. Image-to-Image and Dense Prediction

Image-to-image generation. As shown in Fig. 5, the same checkpoint handles many image-to-image tasks. It can turn edge maps and depth maps into realistic images, colorize grayscale photos, and render photos in anime, Studio Ghibli and watercolor styles. Jeveration achieves strong structural fidelity, with 0.880 HED edge F1, 0.802 depth δ_1 and 0.937 colorization edge F1 on held-out examples (Tab. 2). Within this structure, the text prompt guides the content and appearance. Notably, this control comes at almost no extra cost, since each request is still a single forward pass of 91.7 ms, only 14 ms more than text-to-image generation. These results show that controllable generation is possible in one forward pass, without iterative sampling.

Dense prediction. The same checkpoint is also a strong dense predictor (Fig. 6). As shown in Tab. 3, it reaches 67.7 mIoU on NYUv2, on par with dense-only models trained from the same backbone (67.8 and 67.3), and significantly outperforms DeepLabV3+, with 14.2 higher mIoU and a 32% lower depth error (0.267 versus 0.393 m). Each prediction takes a single forward pass of about 78 ms. Surface normals lag behind, likely because fine local geometry is the most sensitive to sharing capacity across tasks. These results demonstrate that a single checkpoint can serve both image generation and dense prediction.

4.4. Ablation Studies

Shared state and decision training. The shared random state lets all positions follow one image, while different states give different images. Decision training teaches the model to use this state. Without it, the first-phase masked model reaches only 0.32 GenEval in one pass (Fig. 2). With a fixed state instead of a random one, HPSv2.1 drops from 28.27 to 26.63 and GenEval from 0.73 to 0.70 (Tab. 4). The randomness of the shared state thus makes the images more coherent and of higher quality.

Correction target. The correction target brings a further gain in image quality, as removing it lowers HPSv2.1 from

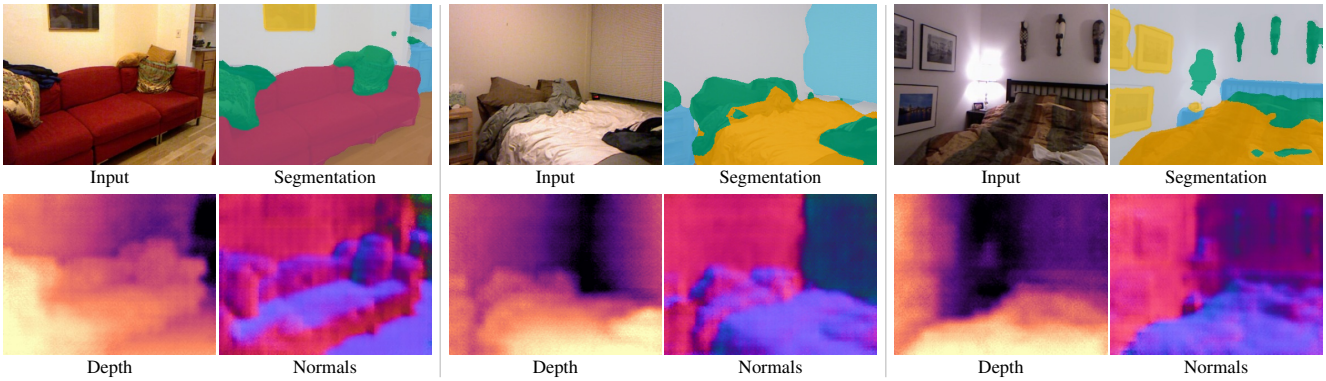


Figure 6. **Dense prediction** of 13-class segmentation, depth and surface normals on NYUv2 test pictures with the same checkpoint.

Table 4. **Ablation studies.** Each block is a separate experiment that changes one component, so values are comparable only within a block. From top to bottom, the blocks start from a text-to-image-only model, an earlier run on all tasks and the first-phase masked model. Highlighted rows are the settings Jeveration uses.

Variant	GenEval↑	HPS↑
w/ random state	**0.73**	**28.27**
w/ fixed state	0.70	26.63
w/ correction target	**0.74**	**28.03**
w/o correction target	0.73	27.92
Image head (wide + codebook)	**0.64**	**25.42**
Image head (wide only)	0.60	25.34
Image head (codebook only)	0.59	25.04

28.03 to 27.92 at similar GenEval (Tab. 4). This suggests that distribution matching already determines the composition, while the correction refines details.

Image-head branches. In the first-phase masked model, the two branches together outperform either branch alone, with 0.64 GenEval against 0.60 and 0.59 for the wide and codebook branches (Tab. 4). We hypothesize that they are complementary, as the codebook branch inherits the structure of the tokenizer's codebook while the wide branch adds capacity beyond this structure.

Source conditioning. Both the generated codes and the source image are essential for image-to-image generation. Even without the source, the decoder recovers most of the input structure from the codes alone (0.792 HED edge F1 and 0.767 depth δ_1), indicating that the VLM transfers the input layout into its decisions. Feeding the source image to the decoder brings further improvements, raising HED edge F1 to 0.880 and colorization edge F1 from 0.757 to 0.937 (Tab. 2). This suggests that the 128 codes capture the global structure, while the source image supplies the fine details.

Table 5. **Joint training across tasks.** Training on all tasks costs almost no text-to-image quality. Both models get about the same amount of text-to-image training, so the difference comes from adding image-to-image and dense tasks.

Training tasks	GenEval↑	HPS↑	DPG↑
Text-to-image only	0.73	**28.28**	**78.61**
All tasks (Jeveration)	**0.74**	28.06	78.24

Joint training across tasks. Joint training extends the model to image-to-image generation and dense prediction at almost no cost to text-to-image quality. As shown in Tab. 5, Jeveration performs on par with a model trained only on text-to-image, with comparable GenEval (0.74 versus 0.73), HPSv2.1 (28.06 versus 28.28) and DPG (78.24 versus 78.61). For dense prediction, it matches dense-only models in segmentation, with only a modest drop in depth and normal accuracy (Tab. 3). This demonstrates that a single set of adapters on a frozen VLM can accommodate diverse visual tasks with little interference.

5. Conclusion

Jeveration introduces a third paradigm for image generation, in which a System One model returns a whole image as one joint decision in a single forward pass. Built on a frozen VLM, it conditions all code positions on a shared random state to encourage global consistency, while a masked model trained beforehand provides the target conditional distributions for parallel code prediction. One checkpoint performs text-to-image generation, image-to-image generation and dense prediction, and joint training across tasks comes at almost no cost to text-to-image quality. Jeveration generates 512-px images in 78 ms on a single H200, achieves strong structural fidelity in image-to-image generation, and delivers segmentation performance comparable to dense-only models. We expect this paradigm to extend to further visual tasks and modalities.

References

[1] Diogo Almeida. Introducing System One Models & Jev. TypeSafe AI Blog, 2026. 2, 3

[2] Carlos André Antunes, Aldry Albuquerque, Gabriel de Melo Schmith, Natã Batista, Pedro Oziel Escobar da Silva, Stefano Benko Rossi, Lingrid Dietrich Lorenz da Silva, Flávia Cristina da Costa, Thiago Bertotto, and Luis Fernando Fausto. Narrow decision models: A system-1 architecture for calibrated, typed decisions. *arXiv preprint arXiv:2610.25553*, 2026. 3

[3] Jinbin Bai, Tian Ye, Wei Chow, Enxin Song, Qing-Guo Chen, Xiangtai Li, Zhen Dong, Lei Zhu, and Shuicheng Yan. Meissonic: Revitalizing masked generative transformers for efficient high-resolution text-to-image synthesis. In *The Thirteenth International Conference on Learning Representations*, 2024. 3

[4] Huanqia Cai, Sihan Cao, Ruoyi Du, Peng Gao, Aiming Hao, Steven Hoi, Zhaohui Hou, Shijie Huang, Dengyang Jiang, Yuming Jiang, et al. Z-image: An efficient image generation foundation model with single-stream diffusion transformer. *arXiv preprint arXiv:2511.22699*, 2025. 6

[5] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11305–11315. IEEE, 2022. 3, 4

[6] Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023. 3, 4

[7] Junsong Chen, Chongjian Ge, Enze Xie, Yue Wu, Lewei Yao, Xiaozhe Ren, Zhongdao Wang, Ping Luo, Huchuan Lu, and Zhenguo Li. PIXART-Σ: Weak-to-strong training of diffusion transformer for 4K text-to-image generation. In *European Conference on Computer Vision*, pages 74–91. Springer, 2024. 6

[8] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. PixArt-α: Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *International conference on learning representations*, pages 57611–57640, 2024. 2, 6

[9] Jiuhai Chen, Zhiyang Xu, Xichen Pan, Yushi Hu, Can Qin, Tom Goldstein, Lifu Huang, Tianyi Zhou, Saining Xie, Silvio Savarese, et al. Blip3-o: A family of fully open unified multimodal models-architecture, training and dataset. *arXiv preprint arXiv:2505.09568*, 2025. 2

[10] Junsong Chen, Shuchen Xue, Yuyang Zhao, Jincheng Yu, Sayak Paul, Junyu Chen, Han Cai, Song Han, and Enze Xie. Sana-sprint: One-step diffusion with continuous-time consistency distillation. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16185–16195. IEEE, 2025. 3, 6

[11] Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. Janus-pro: Unified multimodal understanding and generation with data and model scaling. *arXiv preprint arXiv:2501.17811*, 2025. 2, 6

[12] Zehua Cheng, Wei Dai, and Jiahao Sun. this-that-model-1.0: A typed decision model that decides in 30 ms, for a millionth of a cent. *arXiv preprint arXiv:2609.23886*, 2026. 2, 3

[13] Zhenglin Cheng, Peng Sun, Jianguo Li, and Tao Lin. Twinflow: Realizing one-step generation on large models with self-adversarial flows. In *International Conference on Learning Representations*, pages 151237–151264, 2026. 3

[14] Chaorui Deng, Deyao Zhu, Kunchang Li, Chenhui Gou, Feng Li, Zeyu Wang, Shu Zhong, Weihao Yu, Xiaonan Nie, Ziang Song, et al. Emerging properties in unified multimodal pretraining. *arXiv preprint arXiv:2505.14683*, 2025. 2

[15] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024. 6

[16] Rongyao Fang, Aldrich Yu, Chengqi Duan, Linjiang Huang, Shuai Bai, Yuxuan Cai, Kun Wang, Si Liu, Xihui Liu, and Hongsheng Li. Flux-reason-6m & prism-bench: A million-scale text-to-image reasoning dataset and comprehensive benchmark. In *International Conference on Learning Representations*, pages 63157–63186, 2026. 6

[17] Ying Fang, Zhifu Gao, Qing Gu, Jiajun He, Kangqi Jing, Xiangang Li, Yu Pu, and Wupeng Wang. Qwen-audio-jev: A unified interface for audio decisions. *arXiv preprint arXiv:2610.21614*, 2026. 3

[18] Zhengyang Geng, Mingyang Deng, Xingjian Bai, Zico Kolter, and Kaiming He. Mean flows for one-step generative modeling. *Advances in Neural Information Processing Systems*, 38:75460–75482, 2026. 3

[19] Jonathan Geuter, Andrew Cheng, and David Alvarez-Melis. Releasing the inner jev: A simple recipe for open decision models. *arXiv preprint arXiv:2610.17205*, 2026. 3

[20] Dhruba Ghosh, Hannaneh Hajishirzi, and Ludwig Schmidt. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023. 6

[21] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 2, 4

[22] Xiwei Hu, Rui Wang, Yixiao Fang, Bin Fu, Pei Cheng, and Gang Yu. Ella: Equip diffusion models with llm for enhanced semantic alignment. *arXiv preprint arXiv:2403.05135*, 2024. 6

[23] Hong Huang, Chenhongyi Yang, Junzhe Sun, Animesh Sinha, Wuyang Chen, and Yifan Jiang. Omni-diffusion-distill: Few-step distillation of unified multimodal diffusion large language models. *arXiv preprint arXiv:2610.10990*, 2026. 3

[24] Dongwon Kim, Ju He, Qihang Yu, Chenglin Yang, Xiaohui Shen, Suha Kwak, and Liang-Chieh Chen. Democratizing text-to-image masked generative models with compact text-aware one-dimensional tokens. In *2025 IEEE/CVF In-*

ternational Conference on Computer Vision (ICCV), pages 18442–18452. IEEE, 2025. 4

[25] Black Forest Labs. Flux. https://github.com/black-forest-labs/flux, 2024. 2, 6

[26] Baoteng Li, Wenzhuo Wu, Kongming Liang, and Zhanyu Ma. Visual jev rewards: Reference-bound verification for multi-subject image generation. *arXiv preprint arXiv:2610.09328*, 2026. 2, 3

[27] Lichen Li, Ziping Dong, Sicheng Feng, Qi Li, and Xinchao Wang. So-arena: Benchmarking the capability spectrum and reliability of system one models. *arXiv preprint arXiv:2611.15777*, 2026. 2

[28] Qi Li and Xinchao Wang. Jeva: System one models from first principles. *arXiv preprint arXiv:2610.18773*, 2026. 3

[29] Yi Li, Hongze Shen, Lexiang Tang, Xin Li, Xinpeng Ding, Yinsong Liu, Deqiang Jiang, Xing Sun, and Xiaomeng Li. Densemllm: Standard multimodal llms for dense prediction. *arXiv preprint arXiv:2602.14134*, 2026. 3

[30] Baijiong Lin and Yu Zhang. Libmtl: A python library for deep multi-task learning. *Journal of Machine Learning Research*, 24(209):1–7, 2023. 7

[31] Baijiong Lin, Weisen Jiang, Pengguang Chen, Feiyang Ye, Yu Zhang, Ying-Cong Chen, Shu Liu, Ivor W Tsang, and James T Kwok. Dual-balancing for multi-task learning. *Neural Networks*, page 108317, 2025. 7

[32] Jiazhen Liu and Long Chen. Robust reasoning segmentation for open-world mllms without catastrophic forgetting. *International Journal of Computer Vision*, 134(12):515, 2026. 3

[33] Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. Unified-io 2: Scaling autoregressive multimodal models with vision, language, audio, and action. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26429–26445. IEEE, 2024. 2

[34] Simian Luo, Yiqin Tan, Longbo Huang, Jian Li, and Hang Zhao. Latent consistency models: Synthesizing high-resolution images with few-step inference. *arXiv preprint arXiv:2310.04378*, 2023. 6

[35] Yiyang Ma, Xingchao Liu, Xiaokang Chen, Wen Liu, Chengyue Wu, Zhiyu Wu, Zizheng Pan, Zhenda Xie, Haowei Zhang, Xingkai Yu, et al. Janusflow: Harmonizing autoregression and rectified flow for unified multimodal understanding and generation. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7739–7751. IEEE, 2025. 6

[36] Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Jiuhai Chen, Kunpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025. 2

[37] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182. IEEE, 2023. 4

[38] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *European conference on computer vision*, pages 746–760. Springer, 2012. 6

[39] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023. 3

[40] Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024. 2

[41] Qwen Team. Qwen3.5: Towards native multimodal agents, 2026. 2, 4

[42] TypeSafe AI. Box: TypeSafe documentation. https://docs.typesafe.ai/primitives/box, 2026. 2, 3

[43] Wei Wang, Wenqiao Zhang, Zheqi Lv, Zheng Dong, Yingjie Niu, Jiaqi Zhu, Jun Xiao, Chao Li, and Yueting Zhuang. Jevvla: Jev-inspired energy-based adaptation for vision-language-action policies. *arXiv preprint arXiv:2610.35010*, 2026. 2, 3

[44] Xinlong Wang, Yufeng Cui, Jinsheng Wang, Fan Zhang, Yueze Wang, Xiaosong Zhang, Zhengxiong Luo, Quan Sun, Zhen Li, Yuqi Wang, et al. Multimodal learning with next-token prediction for large multimodal models. *Nature*, 650 (8101):327–333, 2026. 2

[45] Chengyue Wu, Xiaokang Chen, Zhiyu Wu, Yiyang Ma, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, Chong Ruan, et al. Janus: Decoupling visual encoding for unified multimodal understanding and generation. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12966–12977. IEEE, 2025. 2

[46] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025. 6

[47] Xiaoshi Wu, Yiming Hao, Keqiang Sun, Yixiong Chen, Feng Zhu, Rui Zhao, and Hongsheng Li. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023. 6

[48] Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li, Ligeng Zhu, Yao Lu, et al. Sana: Efficient high-resolution text-to-image synthesis with linear diffusion transformers. In *The Thirteenth International Conference on Learning Representations*, 2025. 2, 6

[49] Jinheng Xie, Weijia Mao, Zechen Bai, David Junhao Zhang, Weihao Wang, Kevin Qinghong Lin, Yuchao Gu, Zhijie Chen, Zhenheng Yang, and Mike Zheng Shou. Show-o: One single transformer to unify multimodal understanding and generation. In *International Conference on Learning Representations*, pages 28240–28264, 2025. 2, 6

[50] Chenkai Xu, Xu Wang, Zhenyi Liao, Yishun Li, Tianqi Hou, and Zhijie Deng. Unicm: A unified consistency model for efficient multimodal generation and understanding. In *Findings of the Association for Computational Linguistics: ACL 2026*, pages 18214–18233, 2026. 3, 6

[51] Ling Yang, Ye Tian, Bowen Li, Xinchen Zhang, Ke Shen, Yunhai Tong, and Mengdi Wang. Mmada: Multimodal large

diffusion language models. *Advances in Neural Information Processing Systems*, 38:138867–138907, 2026. 6

[52] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and William T Freeman. Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems*, 37:47455–47487, 2024. 3, 6

[53] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6613–6623. IEEE, 2024. 3

[54] Qihang Yu, Mark Weber, Xueqing Deng, Xiaohui Shen, Daniel Cremers, and Liang-Chieh Chen. An image is worth 32 tokens for reconstruction and generation. *Advances in Neural Information Processing Systems*, 37:128940–128966, 2024. 4

[55] Andrew Zhang, Trinh TL Vuong, Sophia J Wagner, Dongmin Bang, Tong Ding, Yuchang Su, Mary Cipperman, Christina S Chen, Judith Lefkes, Ramin Kahidi, et al. A general-purpose system one decision model for clinical records. *arXiv preprint arXiv:2610.27099*, 2026. 3

[56] Chunting Zhou, Lili Yu, Arun Babu, Kushal Tirumala, Michihiro Yasunaga, Leonid Shamis, Jacob Kahn, Xuezhe Ma, Luke Zettlemoyer, and Omer Levy. Transfusion: Predict the next token and diffuse images with one multi-modal model. In *International Conference on Learning Representations*, pages 6446–6469, 2025. 2

[57] Yuanzhi Zhu, Xi Wang, Stéphane Lathuilière, and Vicky Kalogeiton. Di [m] o: Distilling masked diffusion models into one-step generator. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18606–18618. IEEE, 2025. 3, 4, 5

[58] Yuanzhi Zhu, Xi Wang, Stéphane Lathuilière, and Vicky Kalogeiton. Soft-di [m] o: Improving one-step discrete image generation with soft embeddings. In *International Conference on Learning Representations*, pages 2884–2931, 2026. 3

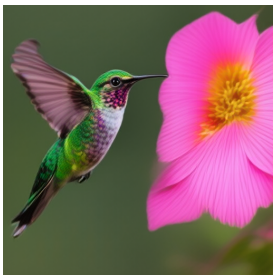
[59] Kai Zou, Zhiwei Zhao, Bin Liu, and Nenghai Yu. Advancing aesthetic image generation via composition transfer. *International Journal of Computer Vision*, 134(5):252, 2026. 6

Jeveration: System One Models Now Generate

Supplementary Material

1. Natural Scenes and Objects

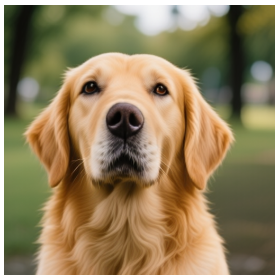
All examples use the final Jeveration checkpoint, with one forward pass per task.



A hummingbird hovering next to a pink hibiscus flower . . .



A portrait of a horse's head with a flowing white mane . . .



A golden retriever's face looking up at the camera . . .



A glowing jellyfish drifting in deep blue ocean water . . .



A turquoise glacier lake surrounded by jagged snowy peaks . . .



Red sand dunes under a starry Milky Way night sky.



A lighthouse on a rocky cliff at sunset . . .



An ancient stone temple in a jungle overgrown with vines.



A stack of fluffy pancakes with blueberries and maple syrup . . .



A slice of strawberry shortcake with whipped cream . . .



A bowl of ramen with a soft-boiled egg . . . and rising steam . . .



A golden croissant . . . next to a cup of coffee . . .



A glass terrarium with succulents and moss on a sunny shelf.



A bonsai tree in a ceramic pot . . .



A stack of old books with a cup of tea and autumn leaves . . .



A minimalist modern living room . . . overlooking the ocean.

Figure 1. **Additional text-to-image results.** Selected 512-px outputs across animals, landscapes, food and everyday scenes, with excerpts of their prompts.

2. Artistic Styles and Portraits



A Japanese ukiyo-e woodblock print of a great wave …



A unicorn head made of clear crystal glass …



A forest spirit with antlers made of branches and glowing moss …



A claymation-style tiny village on a green hill …



An oil painting of a stormy sea with a sailing ship …



A Chinese ink wash painting of misty mountains …



An impressionist oil painting of a bouquet of peonies …



Layered paper-cut art of a forest and mountains …



An anime-style countryside landscape …



A gouache illustration of a fox in a snowy winter forest …



A dreamy pastel watercolor of a hot air balloon …



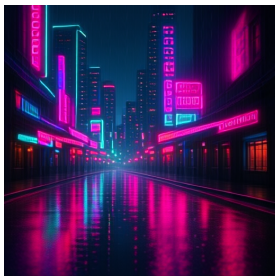
A paper quilling artwork of a hummingbird and flowers …



A close-up portrait of a woman with glowing golden makeup …



A studio portrait of a girl with colorful butterfly face paint.



A neon-lit futuristic city at night with rain reflections …



A smiling young man with curly hair in soft natural light …

Figure 2. **Fantasy, artistic styles and portraits.** Additional selected outputs of the same single-pass checkpoint.

3. Image-to-Image Generation and Dense Prediction

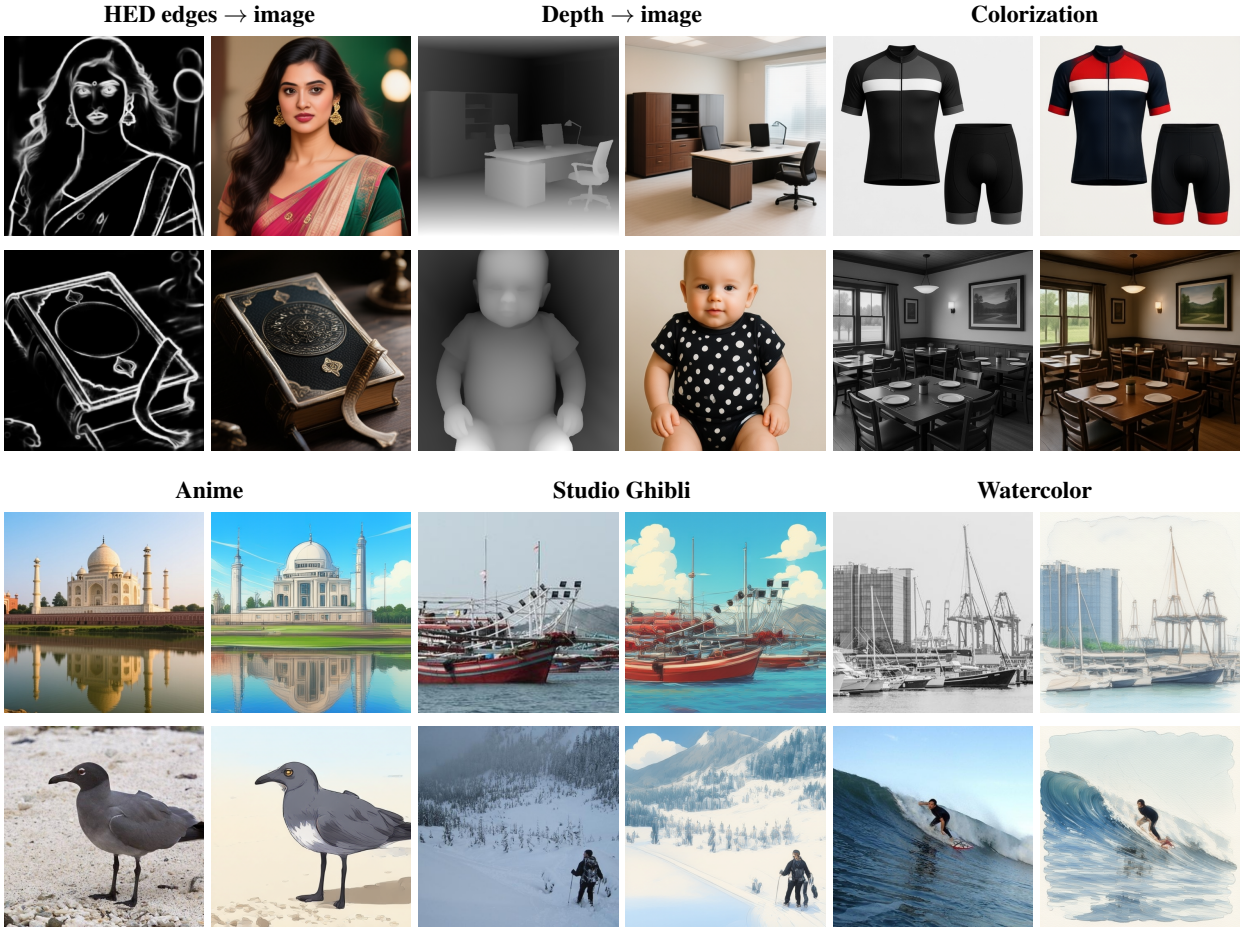


Figure 3. **Additional image-to-image results.** Each pair shows the condition or source image on the left and Jeveration's output on the right. Structure tasks also receive a caption, while style tasks use a fixed instruction for each style.

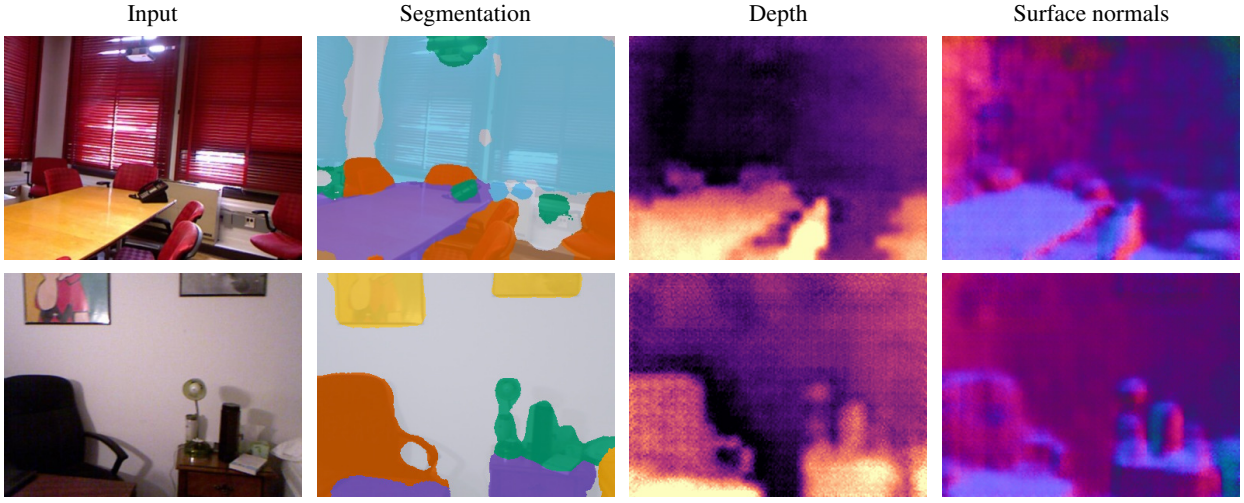


Figure 4. **Dense outputs from the same checkpoint.** NYUv2 test images and single-pass predictions, arranged by task for comparison. Nearer depth is brighter.

4. Implementation and Evaluation

Model. Jeveration keeps the Qwen3.5-4B backbone and its vision tower frozen. All tasks share rank-256 LoRA adapters, and task-specific heads read the decision tokens. Image generation uses 128 TA-TiTok codes from an 8,192-entry codebook, and dense prediction uses 432 spatial queries. About 1.0B parameters are trainable in total, including the adapters, input modules and output heads.

Training. Training has two phases, masked prediction followed by decision training with distribution matching and correction (Tab. 1). We use AdamW with an EMA of 0.999 on four NVIDIA RTX PRO 6000 GPUs. The correction re-masks the 50% least confident positions of a sample, which the frozen first-phase model completes in four rounds at CFG 3, with $\lambda = 1$.

Table 1. **Main training settings.** The backbone remains frozen in every stage.

Training stage	Steps	LR	Batch
Phase 1, generation only	80k	2×10^{-4}	256
Phase 1, all tasks	12k	5×10^{-5}	256
Phase 2, distribution matching	5k	5×10^{-6}	64
Phase 2, with correction	1.7k	5×10^{-6}	64

Decoding. We adapt the text-aware decoder of TA-TiTok to 512 px and train it to reconstruct images from their codes. A source-conditioning branch is added for image-to-image tasks. The decoder relies on both of its inputs (Tab. 2). Replacing the codes with those of another example lowers HED edge F1 from 0.896 to 0.552 and depth δ_1 from 0.825 to 0.346, while replacing the source image lowers them to 0.652 and 0.767. This shows that the codes carry most of the structure and the source image refines it.

Table 2. **Code and source swap.** Each row replaces one decoder input with that of another example, on a shorter training run (32 examples per structure task, 223 style examples).

Decoder input	HED↑	Depth↑	Style PSNR↑
Correct codes and source	**0.896**	**0.825**	**15.43**
Wrong codes, correct source	0.552	0.346	9.84
Correct codes, wrong source	0.652	0.767	14.71

Evaluation. For text-to-image generation, GenEval uses its 553 prompts with four samples each, HPSv2.1 uses 3,200 prompts, and DPG-Bench scores 1,064 prompts as 2×2 grids. For image-to-image generation, structure fidelity is measured on 341 held-out examples per task. We extract the structure from each output and compare it with that of the reference image, using HED edge F1, depth δ_1 after scale-and-shift alignment, and Canny edge F1 for colorization. Style PSNR compares each output with its paired target on 223 held-out examples. For dense prediction, we report the standard NYUv2 metrics on the 654 test images. The dense-only models use the same backbone and initialisation as Jeveration. All latencies are end-to-end times from prompt to decoded image on one H200 at batch 1, reported as the median of five calls after two warm-up calls.