

# Symmetrical Flow Matching: Unified Image Generation, Segmentation, and Classification with Score-Based Generative Models

Francisco Caetano<sup>1</sup>, Christiaan Viviers<sup>1</sup>, Peter H.N. De With<sup>1</sup>, Fons van der Sommen<sup>1</sup>

<sup>1</sup>Eindhoven University of Technology, The Netherlands  
f.caetano@tue.nl

## Abstract

Flow Matching has emerged as a powerful framework for learning continuous transformations between distributions, enabling high-fidelity generative modeling. This work introduces Symmetrical Flow Matching (SymmFlow), a new formulation that unifies semantic segmentation, classification, and image generation within a single model. Using a symmetric learning objective, SymmFlow models forward and reverse transformations jointly, ensuring bi-directional consistency, while preserving sufficient entropy for generative diversity. A new training objective is introduced to explicitly retain semantic information across flows, featuring efficient sampling while preserving semantic structure, allowing for one-step segmentation and classification without iterative refinement. Unlike previous approaches that impose strict one-to-one mapping between masks and images, SymmFlow generalizes to flexible conditioning, supporting both pixel-level and image-level class labels. Experimental results on various benchmarks demonstrate that SymmFlow achieves state-of-the-art performance on semantic image synthesis, obtaining FID scores of 11.9 on CelebAMask-HQ and 7.0 on COCO-Stuff with only 25 inference steps. Additionally, it delivers competitive results on semantic segmentation and shows promising capabilities in classification tasks.

**Code** — <https://github.com/caetas/SymmetricFlow>

## 1 Introduction

Comprehending semantic content is a key challenge in computer vision. Classification (Krizhevsky, Sutskever, and Hinton 2012; He et al. 2016; Dosovitskiy et al. 2020) and segmentation (Ronneberger, Fischer, and Brox 2015; Xie et al. 2021; Cheng et al. 2022) allow models to analyze and structure images, while generative modeling enables the synthesis of new content (Radford, Metz, and Chintala 2015; Ho, Jain, and Abbeel 2020; Rombach et al. 2022). Ideally, a unified framework would bridge these tasks, allowing models to both interpret and generate images in a two-way manner. It is conceivable that the ability to accurately comprehend and disentangle visual structures facilitates the generation of more semantically coherent and visually realistic images.

Conversely, strong generative capabilities may aid in learning more expressive representations of images, as generating plausible content requires an implicit understanding of object relationships, textures, and context (He et al. 2022). These forward and backward relationships suggest that advances in one direction could naturally benefit the other, which fuels the pursuit of models that integrate both understanding and synthesis within a single cohesive framework.

Most existing vision-only approaches treat these tasks individually. For classification, the extracted features are used in fully connected layers to obtain probability predictions for each class, while for segmentation, the features are used to train decoders for dense predictions. On the other hand, generative models, such as generative adversarial networks (GANs) (Karras et al. 2021), diffusion models (Karras et al. 2022), score matching models (Song et al. 2020), and Flow Matching models (Lipman et al. 2022), synthesize images from a prior distribution. Although recent work has explored diffusion models for classification (Li et al. 2023a) and segmentation (Wu et al. 2024), these adaptations introduce significant limitations: classification is slow due to the need for iterative sampling across all possible classes, and segmentation frameworks are restricted to generating masks, lacking the ability to map back to realistic images.

Recent works, such as SemFlow (Wang et al. 2024b) and DepthFM (Gui et al. 2024), have sought to unify image generation with semantic segmentation and depth estimation within a single generative framework. However, these models still suffer from the following key limitations: (1) they do not perform classification, restricting their applicability; (2) the image quality remains inferior to that of purely generative models; and (3) the models enforce a strict one-to-one mapping between segmentation or depth masks and images, requiring them to have the same number of channels, which limits their flexibility.

To address these limitations, we propose Symmetrical Flow Matching (SymmFlow), a novel Flow Matching training objective that enforces dual symmetrical sampling, enabling segmentation, classification, and image synthesis to be conditioned in a mutual two-sided relationship. This concept leads to the following specific contributions. (a) SymmFlow unifies segmentation and classification within a single framework, performing both tasks in fewer steps, while retaining the ability to generate high-quality im-

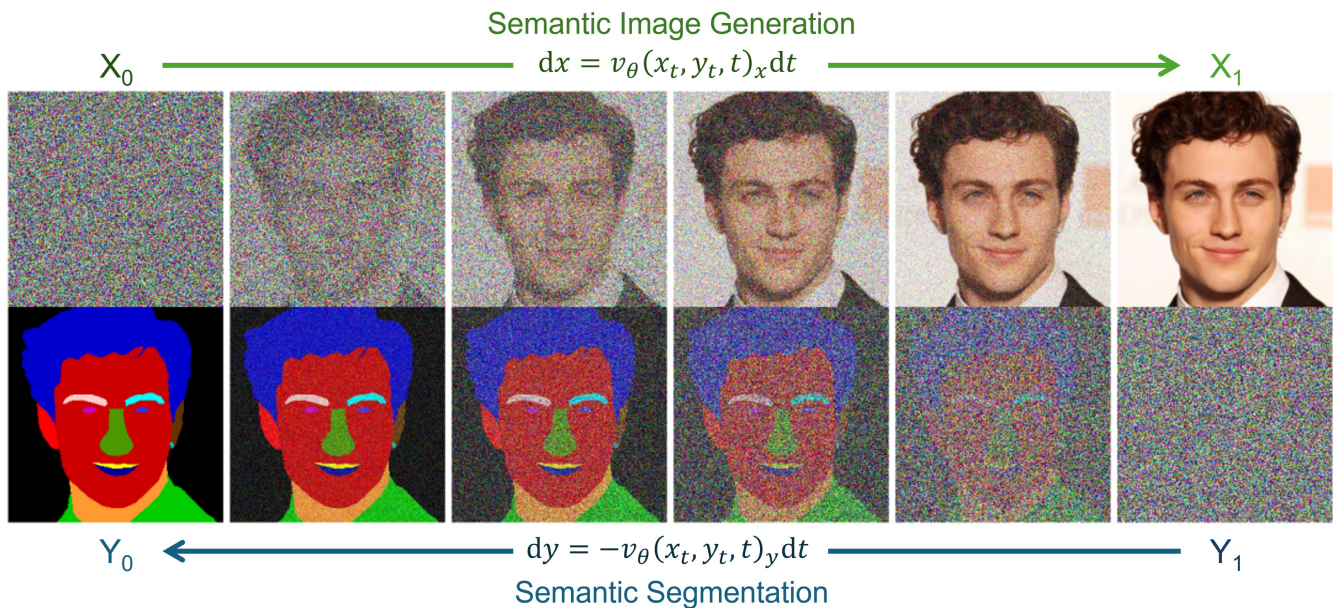


Figure 1: Symmetrical Flow Matching jointly models semantic segmentation and generation as opposing flows. Noise transitions into an image while a label evolves into noise and vice versa. This symmetry maintains entropy for generation while enforcing semantic consistency. Image Y can represent semantic content of any type, from dense masks to global labels, enabling applications like classification and segmentation.

ages through Flow Matching. (b) SymmFlow enhances image synthesis quality over prior methods by leveraging the bi-directionality of Flow Matching. (c) SymmFlow alleviates the strict one-to-one channel constraint between segmentation masks and images, allowing greater flexibility in conditioning and generalization. We validate SymmFlow on toy problems as well as standard benchmarks for classification, segmentation, and image synthesis, demonstrating its effectiveness as a unified model for both discriminative and generative tasks.

## 2 Related Work

### 2.1 Flow Matching

Flow Matching (FM) (Lipman et al. 2022; Liu, Gong, and Liu 2022) is a generative modeling framework that learns a velocity field to transform a source distribution into a target distribution through a continuous flow. By parameterizing this field with a neural network, the model becomes a Neural ODE (Chen et al. 2018), allowing efficient sampling via numerical integration. It has advanced state-of-the-art performance in diverse applications, including image (Esser et al. 2024) and audio generation (Vyas et al. 2023), as well as protein modeling (Huguet et al. 2024) and robotics (Black et al. 2024). FM generalizes Continuous Normalizing Flows (CNFs) (Chen et al. 2018; Grathwohl et al. 2018) by eliminating the need for simulation during training, thereby making it computationally efficient. It also provides a unifying perspective on generative modeling, encompassing diffusion models, which can be seen as a special case where the probability path is defined via stochastic differential equations (Lipman et al. 2024).

### 2.2 Generative Classifiers

Several seminal works (Hinton 2007; Ranzato et al. 2011) have emphasized the importance of modeling the data distribution to enhance discriminative feature learning. Early approaches trained deep belief networks (Hinton, Osindero, and Teh 2006) to encode image data as latent representations, which were then used for recognition tasks. More recent advances in generative modeling have demonstrated the ability to learn efficient representations for global prediction tasks (He et al. 2022). In addition, generative models have been shown to improve adversarial robustness and calibration (Huang et al. 2020). However, most prior work either jointly trains generative and discriminative models, or fine-tunes generative representations for downstream tasks. Diffusion Classifier (Li et al. 2023a) specifically investigates the effectiveness of using diffusion models as image classifiers, albeit with severe inference constraints.

### 2.3 Semantic Segmentation

Semantic segmentation aims to assign a semantic label to each pixel in an image. Conventional approaches rely on discriminative models, combining a strong feature extraction backbone with a task-specific decoder head for mask prediction (Cheng, Schwing, and Kirillov 2021; Cheng et al. 2022; Ding et al. 2023). Recent work has explored diffusion models for segmentation (Amit et al. 2021; Wang et al. 2024a; Gu, Chen, and Xu 2024), typically leveraging them as feature extractors within a discriminative framework. A key motivation behind SemFlow is the interpretation that diffusion models struggle to align their stochastic nature with the deterministic requirements of semantic segmenta-

tion. For this reason, SemFlow introduces rectified flows. In contrast, SymmFlow embraces the probabilistic nature of segmentation, accounting for inter-observer variability.

## 2.4 Semantic Image Generation

Semantic image generation, the inverse of semantic segmentation, focuses on generating realistic images from semantic layouts (Li, Zhang, and Malik 2019; Liu et al. 2019; Lv et al. 2022). Existing approaches generally fall into two categories. (1) GAN-based models (Zhu et al. 2017; Isola et al. 2017; Wang et al. 2018), although many of these methods struggle with mode collapse and produce only unimodal outputs. (2) Diffusion models, which treat semantic generation as a conditional generation task where semantic masks act as control signals (Wang et al. 2022). Some methods further integrate additional conditioning to enhance coherence (Zhang, Rao, and Agrawala 2023). However, these approaches often adopt asymmetric architectures with unidirectional generators, thereby complicating the unification of semantic segmentation and image generation.

## 3 Symmetrical Flow Matching

Symmetrical Flow Matching (SymmFlow) unifies semantic segmentation and semantic synthesis as opposing flow processes, as illustrated in Figure 1. Given a data distribution  $X$  (e.g. images) and a semantic representation  $Y$  (e.g. masks or class labels), SymmFlow models bi-directional flows between them. The forward process transforms  $X$  from noise, while simultaneously evolving  $Y$  towards a noise-corrupted state. The reverse process inverts these transitions, allowing for the generation of  $Y$  from  $X$ . Crucially,  $Y$  is not restricted to having the same dimensionality as  $X$ , enabling flexible conditioning, such as global class labels for classification. This symmetrical formulation ensures sufficient entropy for image generation while preserving the semantic structure, making SymmFlow a generalizable framework for both segmentation and synthesis. The training procedure and objective are formalized in Section 3.1. Techniques for obtaining segmentation and classification predictions using the proposed SymmFlow model are discussed in Section 3.2. The importance of label dequantization for stable training is examined in Section 3.3, and the complete framework is validated using a synthetic example presented in Section 3.4.

### 3.1 Training Objective

Symmetrical Flow Matching jointly models semantic segmentation and synthesis as opposing flows, enabling bi-directional transformations between images and semantic content. The model learns a velocity field that transports  $X$  from noise ( $X_0$ ), while simultaneously evolving  $Y$  into noise and vice versa. For each sample, a time variable  $t$  is extracted from  $\mathcal{U}(0, 1)$ , and the inputs are perturbed via a convex combination with Gaussian noise. As a result, the perturbed samples  $x_t$  (forward) and  $y_t$  (backward) are specified by

$$\begin{aligned} x_t &= (1-t)\xi_x + tx, \\ y_t &= (1-t)y + t\xi_y, \end{aligned} \quad (1)$$

where  $\xi_x, \xi_y$  are independent noise terms extracted from  $\mathcal{N}(0, I)$ . The optimal transport velocity fields are given by

$$\begin{aligned} v_x &= x - \xi_x, \\ v_y &= \xi_y - y, \\ v &= (v_x, v_y), \end{aligned} \quad (2)$$

which describe the ideal directions to reverse the perturbation. Figure 2 illustrates the optimal transport approach in Symmetrical Flow Matching, showing the transformation between data distributions and the Gaussian intermediary.




Figure 2: Illustration of the optimal transport between the data distributions  $X$  and  $Y$ , and the intermediate Gaussian distribution.

The model  $v_\theta(x_t, y_t, t)$  is trained to jointly approximate both flows by minimizing the squared error, specified by

$$\mathcal{L} = \mathbb{E}_{x, y, t} [\|v_\theta(x_t, y_t, t) - v\|^2]. \quad (3)$$

### 3.2 Classification and Segmentation

A common approach to classification using conditional generative models relies on Bayes' theorem to compute the posterior probability of a class  $c$  given an input image  $X$ . Given a generative model that learns the conditional distribution  $p_\theta(x | c)$ , classification is performed as

$$p_\theta(c_i | x) = \frac{p(c_i)p_\theta(x | c_i)}{\sum_j p(c_j)p_\theta(x | c_j)}. \quad (4)$$

With a uniform prior over classes, i.e.  $p(c_i) = \frac{1}{N}$ , the prior terms cancel, simplifying to

$$p_\theta(c_i | x) \propto p_\theta(x | c_i). \quad (5)$$

For diffusion models, computing  $p_\theta(x | c)$  is intractable, so an ELBO approximation is used to estimate the posterior distribution

$$p_\theta(c_i | x) = \frac{\exp\{-\mathbb{E}_{t, \epsilon}[\|\epsilon - \epsilon_\theta(x_t, c_i)\|^2]\}}{\sum_j \exp\{-\mathbb{E}_{t, \epsilon}[\|\epsilon - \epsilon_\theta(x_t, c_j)\|^2]\}}. \quad (6)$$

Monte Carlo sampling approximates the expectation by

$$\frac{1}{N} \sum_{i=1}^N \|\epsilon_i - \epsilon_\theta(\sqrt{\alpha_{t_i}}x + \sqrt{1 - \alpha_{t_i}}\epsilon_i, c_j)\|^2. \quad (7)$$

Ultimately, this approach extracts a classifier by evaluating the error between noise predictions for each class.

**Proposed Approach** In contrast to the conventional generative classifier approach, SymmFlow learns a velocity field that transports an input image toward a noise distribution, and vice versa. Classification is performed by integrating the predicted velocity field in an off-the-shelf Ordinary Differential Equation (ODE) solver

$$y_0 = y_1 + \int_1^0 v_\theta(x_t, y_t, t)_y dt. \quad (8)$$

This eliminates the need for repeated evaluations across all possible class embeddings, significantly reducing inference time and computational cost. Additionally, the same process can be used to predict segmentation masks. The predicted class is determined as the closest label to the average of the model’s predictions. For segmentation, the class of each pixel is assigned based on the closest predefined class RGB code to the predicted pixel RGB value. This mapping is further explained in Section 3.4.

### 3.3 Dequantization

In line with prior work, we dequantize the labels (classification and segmentation) to a continuous distribution. Dequantization is often used in Normalizing Flows to enhance stability in density modeling. Without it, excessively high likelihoods (dirac deltas) are assigned to a few specific values, causing the model to collapse. A standard approach to dequantization involves adding finite perturbations to the signal to prevent low-entropy distributions from hindering modeling quality. We adopt a similar strategy by applying controlled noise to the class labels  $Y$ , ensuring smoother optimization and preventing degenerated solutions. Specifically, given a discrete label  $Y$ , we define the dequantized representation as

$$Y' = Y + \epsilon, \quad \epsilon \sim U(-\beta, +\beta), \quad (9)$$

where  $U(-\beta, +\beta)$  is a uniform noise term ensuring that the semantic label remains well-defined. This adjustment is crucial for maintaining stability in the reverse flow process.

For models performing classification, we further normalize the label representations to the interval  $[-1, +1]$ , based on their indices, prior to supplying them to the model. The dequantized values then serve as a continuous mask, providing a structured conditioning mechanism for the model.

### 3.4 Toy Example

To illustrate the principles of Symmetrical Flow Matching, we consider a toy example where two classes form the interleaved, nonlinear structures in Figure 3a.

We train a multilayer perceptron (MLP) to model the joint evolution of data points  $X$  and their class representation  $Y$  under Symmetrical Flow Matching. The input to the model consists of the point coordinates  $X$ , a quantized encoding of their class  $Y$  - where Class A follows a uniform distribution in  $[-1.5, -0.5]$  and Class B in  $[+0.5, +1.5]$  — and a time variable  $t$  ranging from 0 to 1. The learned flow is then used to sample trajectories via an Euler ODE solver with 20 timesteps from  $t = 0$  to  $t = 1$ , reconstructing the underlying structure of the data. The resulting distributions,

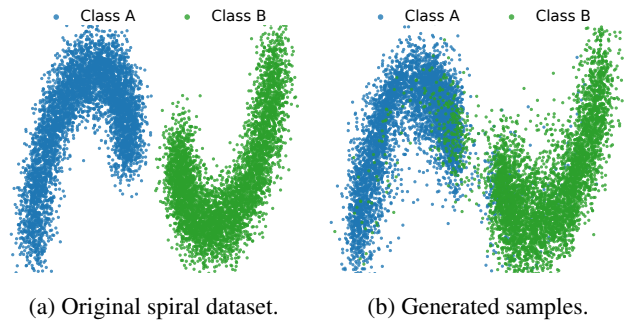


Figure 3: Visualization of the spiral dataset and samples generated by our model.

shown in Figure 3b, demonstrate the ability of our framework to model structured transformations while maintaining class separability.

We also model the reverse process, integrating the learned flow backward from  $t = 1$  to  $t = 0$  to recover class labels. This formulation treats classification as an inverse problem, requiring the model to separate class information as the distribution regresses to its original state. Table 1 shows that the classification accuracy is highest when using a single integration step. This result could be expected: as  $X$  evolves toward a Gaussian distribution, class boundaries blur, making it increasingly difficult for the reverse flow to correctly infer the original labels.

$N_{\text{steps}}$	1	2	5	10	20	50
<b>Acc. (%)</b>	<b>100.0</b>	92.0	87.0	83.6	82.6	82.0

Table 1: Classification accuracy at different numbers of steps using the reverse process of the model.

## 4 Experiments

To evaluate the effectiveness of the proposed SymmFlow model as a unified architecture for semantic segmentation/classification and image synthesis, we perform the following experiments.

### 4.1 Datasets

*Semantic Segmentation and Generation:* The model is evaluated using COCO-Stuff (Caesar, Uijlings, and Ferrari 2018) and CelebAMask-HQ (Lee et al. 2020), which contain 171 and 19 classes, respectively. Images and semantic masks are resized and cropped into  $512 \times 512$  pixels. A detailed analysis of the datasets is provided in Appendix A. *Classification:* We evaluate our model on MNIST (Deng 2012) and CIFAR-10 (Krizhevsky, Nair, and Hinton 2010) for classification, leveraging these low-resolution datasets to assess fundamental capabilities.

### 4.2 Metrics

*Semantic Segmentation and Generation:* For semantic segmentation, we evaluate with mean intersection over

Category	Method	Steps	SS (mIoU $\uparrow$ )		SIS (FID $\downarrow$ / LPIPS $\uparrow$ )	
			CelebAMask-HQ	COCO-stuff	CelebAMask-HQ	COCO-stuff
SS	DML-CSR (Zheng et al. 2022)	1	77.8	—	$\times$	$\times$
	SegFace (Narayan, Vs, and Patel 2025)	1	<b>81.6</b>	—	$\times$	$\times$
	DeeplabV2 (Caesar, Uijlings, and Ferrari 2018)	1	—	33.2	$\times$	$\times$
	MaskFormer (Cheng, Schwing, and Kirillov 2021)	1	—	37.1	$\times$	$\times$
	SegFormer (Xie et al. 2021)	1	—	<b>46.7</b>	$\times$	$\times$
SIS	pix2pixHD (Wang et al. 2018)	1	$\times$	$\times$	54.7 / 0.529	111.5 / —
	SPADE (Park et al. 2019)	1	$\times$	$\times$	42.2 / 0.487	33.9 / —
	SC-GAN (Wang et al. 2021)	1	$\times$	$\times$	19.2 / 0.395	18.1 / —
	BBDM (Li et al. 2023b)	200	$\times$	$\times$	21.4 / 0.370	—
	ControlNet (Zhang, Rao, and Agrawala 2023)	20	$\times$	$\times$	24.0 / 0.528	36.6 / 0.671
	SDM (Wang et al. 2022)	1000	$\times$	$\times$	18.8 / 0.422	15.9 / 0.518
	SCDM (Ko et al. 2024)	250	$\times$	$\times$	17.4 / 0.418	15.3 / 0.519
SCP-Diff (Gao et al. 2024)	800	$\times$	$\times$	—	11.3 / —	
Both	SemFlow (Wang et al. 2024b)	25	69.4*	35.7*	32.6 / 0.393	90.0 / 0.685*
	SymmFlow (Proposed)	25	69.3	39.6	<b>11.9</b> / 0.464	<b>7.0</b> / 0.609

Table 2: Performance comparison of benchmark solutions on semantic segmentation (SS) and semantic image synthesis (SIS) tasks across both COCO-Stuff and CelebAMask-HQ datasets. The number of steps indicates the number of functions applied to obtain the results. Legend: \* Recomputed by us; — Not available;  $\times$  Not applicable (method cannot perform the task).

union (mIoU). For semantic image synthesis, we assess with the Fréchet inception distance (FID) and learned perceptual image patch similarity (LPIPS). *Classification*: The model is evaluated using classification accuracy. Additional information is available in Appendix D.

### 4.3 Qualitative Evaluation

To complement the quantitative analysis, qualitative results are presented for classification, semantic segmentation, and image generation. For classification, generated samples from MNIST and CIFAR-10 are visualized in Appendix G to evaluate the diversity and fidelity of the synthesized outputs. For semantic segmentation, predicted masks on CelebAMask-HQ and COCO-stuff are compared against ground-truth annotations. For image generation, representative samples synthesized from CelebAMask-HQ and COCO-stuff are provided to illustrate visual quality.

### 4.4 Impact of Inference Steps

The effect of the number of inference steps is evaluated for both tasks. By varying the number of steps, the impact on classification accuracy and segmentation quality is assessed, providing insight into the trade-off between computational efficiency and performance.

### 4.5 Baseline Comparison

*Semantic Segmentation and Generation*: For semantic segmentation on CelebAMask-HQ, DML-CSR (Zheng et al. 2022) and SegFace (Narayan, Vs, and Patel 2025) are adopted as baselines. On COCO-Stuff, DeeplabV2 (Caesar, Uijlings, and Ferrari 2018), MaskFormer (Cheng, Schwing, and Kirillov 2021), and SegFormer (Xie et al. 2021) serve as representative methods. For semantic image synthesis, the evaluation includes pix2pixHD (Wang et al. 2018),

SPADE (Park et al. 2019), SC-GAN (Wang et al. 2021), BBDM (Li et al. 2023b), ControlNet (Zhang, Rao, and Agrawala 2023), SDM (Wang et al. 2022), SCDM (Ko et al. 2024), and SCP-Diff (Gao et al. 2024). SemFlow (Wang et al. 2024b), a prior model aiming to unify both segmentation and synthesis, is also evaluated. Since official checkpoints and semantic segmentation evaluation scripts for SemFlow were not publicly available, the model was retrained and results were recomputed. The scripts are included in SymmFlow’s repository for reproducibility. *Classification*: For the classification task on CIFAR-10, the proposed model is compared to the Diffusion Classifier (Li et al. 2023a). Additionally, the quality of the generated images is assessed by benchmarking against a standard FM setup.

### 4.6 Implementation Details

For pixel-level implementations (e.g., classification), the U-Net architecture introduced in Guided Diffusion (Dhariwal and Nichol 2021) is employed. For latent-space models, the pre-trained VAE from Stable Diffusion (Rombach et al. 2022) is used for image encoding and decoding, alongside the U-Net backbone from Stable Diffusion 2.1. To ensure compatibility with SymmFlow, the number of input channels in the first layer and output channels in the final layer of the U-Net are doubled. Further details on hyperparameters and computational resources are provided in Appendix B.

## 5 Results & Discussion

### 5.1 Semantic Segmentation and Generation

Table 2 reports semantic segmentation and image synthesis results across benchmarks. SymmFlow consistently outperforms prior models in the synthesis task, achieving the lowest FID scores across datasets. The visualizations in Figure 4 confirm that the model produces mask-consistent high-



Figure 4: Non-curated samples generated by the model trained on CelebAMask-HQ (left) and COCO-stuff (right). The top row shows the semantic mask used to condition the model. The bottom row shows the samples after 25 integration steps with the Euler ODE solver.

fidelity samples on CelebAMask-HQ and COCO-Stuff, capturing structural details with strong adherence to the conditioning masks. Although LPIPS is commonly used to assess perceptual diversity, it can be misleading in isolation, as higher LPIPS values may result from poor image quality rather than true variability. Conversely, low LPIPS may indicate mode collapse or data leakage rather than accurate diversity. Therefore, interpreting LPIPS jointly with FID is essential. SymmFlow achieves the best trade-off between image quality and diversity, reflecting strong generative capability without sacrificing semantic alignment.

In the segmentation task, SymmFlow achieves competitive performance compared to specialized segmentation baselines, particularly on COCO-Stuff. As shown in Figure 5, the model demonstrates semantic understanding beyond the ground-truth annotations; for instance, correctly identifying a laptop absent from the label map. However, performance is limited by the low-resolution latent representation ( $64 \times 64 \times 4$ ), which hampers fine-grained accuracy. This is especially evident for small-area classes such as earrings or partially occluded features like eyebrows and ears, where segmentation quality deteriorates due to insufficient spatial detail. Despite this, the global segmentation structure remains coherent. SemFlow, while also achieving strong performance on the segmentation task, falls short on semantic image synthesis. Quantitative results indicate limited visual fidelity, underscoring the limitations of the architecture in jointly modeling segmentation and generation without further design or training adjustments.

SymmFlow achieves this performance using only 25 function evaluations, whereas most diffusion-based methods require hundreds of denoising steps, resulting in significantly improved inference efficiency, as shown in Appendix C.

## 5.2 Classification

Table 3 presents the classification performance comparison. With a single inference step, SymmFlow achieves accuracy comparable to the Diffusion Classifier while being significantly more efficient. By increasing the number of steps to just 25, which is 100 times fewer than required by the Diffusion Classifier, SymmFlow convincingly outperforms it on CIFAR-10. This highlights the efficiency of the approach,

reducing inference time without sacrificing accuracy. Additionally, the model currently uses a simple conditioning strategy based on grayscale intensities, suggesting further improvements can be achieved with more advanced conditioning mechanisms. In Appendix G, Figures 6 and 7 present non-curated samples generated by SymmFlow on MNIST and CIFAR-10, illustrating the model’s ability to produce diverse and visually coherent outputs. These findings indicate that SymmFlow continues to be a robust image generator, enabling efficient conditional control.

Method	Steps	MNIST	CIFAR-10
Diffusion Classifier (Li et al. 2023a)	2,750	—	88.5
SymmFlow (Proposed)	1	99.3	88.2
	25	<b>99.6</b>	<b>90.6</b>

Table 3: Comparison between Diffusion Classifier and SymmFlow on image classification tasks.

## 5.3 Evaluation of Inference Steps

Table 4 evaluates the effect of the number of inference steps on semantic image generation performance across both CelebAMask-HQ and COCO-Stuff. On both datasets, FID and LPIPS scores steadily decrease as the number of steps increases, indicating improvements in image fidelity and structural coherence. In both datasets, the LPIPS trend is particularly revealing: at low step counts, the high LPIPS values do not primarily reflect meaningful diversity, but rather poor visual quality and weak alignment with the conditioning masks. As the number of steps increases, the model produces more coherent and semantically accurate outputs, leading to a decrease in LPIPS. Importantly, the final LPIPS values remain relatively high, but this now reflects genuine variability across samples rather than noise or misalignment, indicating that the model preserves diversity while improving structural fidelity. Appendix E analyzes the impact of the dequantization factors, while Appendix G illustrates the qualitative effects of different numbers of sampling steps.

Table 5 shows that one-step segmentation on CelebAMask-HQ already yields a respectable 65.3




Figure 5: Non-curated segmentation masks generated by the model trained on CelebAMask-HQ (left) and COCO-stuff (right). The top row shows the ground-truth segmentation mask. The middle row shows the image used to condition the model. The bottom row shows the segmentations after 25 integration steps with the Euler ODE solver.

Dataset	Metric	1	2	5	10	20	25
CelebA	FID↓	88.5	73.2	49.5	28.2	14.1	<b>11.9</b>
	LPIPS↓	0.598	0.572	0.522	0.486	0.466	<b>0.464</b>
COCO	FID↓	102.6	83.6	44.3	18.2	8.1	<b>7.0</b>
	LPIPS↓	0.777	0.758	0.704	0.652	0.616	<b>0.609</b>

Table 4: Semantic image generation performance for different numbers of steps on CelebAMask-HQ and COCO-Stuff.

mIoU, rising to its maximum of 70.3 mIoU by two steps before plateauing. On COCO-Stuff, segmentation quality reaches a solid 38.1 mIoU by five steps and continues to improve modestly, peaking at 40.1 mIoU by twenty steps. While image generation clearly benefits from additional inference steps, these results suggest that, when the sole objective is semantic segmentation, far fewer steps can be employed without sacrificing much accuracy.

Dataset	Metric	1	2	5	10	20	25
CelebA	mIoU↑	65.3	<b>70.3</b>	<b>70.3</b>	69.8	69.4	69.3
COCO	mIoU↑	29.3	33.8	38.1	38.9	<b>40.1</b>	39.6

Table 5: Semantic segmentation performance at different numbers of steps on CelebAMask-HQ and COCO-Stuff.

The impact of the inference steps in classification is covered in Appendix F, and summarized in Table 11.

## 6 Limitations & Future Work

One limitation of the current approach is the overall model size. Although SymmFlow uses far fewer inference steps than typical diffusion models, it relies on a large pre-trained Stable Diffusion U-Net backbone. Reducing the computational burden by distilling the model into a one-step variant is a natural next step. Moreover, a one-step formula-

tion could facilitate the integration of more expressive conditioning mechanisms, enabling effective text-based control and the use of more powerful architectures such as MMDiT (Esser et al. 2024). Additionally, fine-tuning the VAE decoder to better align with semantic masks could improve segmentation accuracy on fine-grained details, such as small or occluded regions.

Future work includes extending classification evaluation from the current proof-of-concept stage to datasets such as Food-101 (Bossard, Guillaumin, and Van Gool 2014), ImageNet-1K (Deng et al. 2009), and ObjectNet (Barbu et al. 2019), while also refining the semantic label encoding strategy to improve conditioning. Furthermore, evaluating the model performance on depth estimation tasks, similar to DepthFM, would further demonstrate its versatility. Beyond segmentation, the bi-directionality of the model presents opportunities for applications such as image editing.

## 7 Conclusions

This work introduces Symmetrical Flow Matching, a unified framework that models segmentation, classification, and image generation as opposing flows within a single architecture. Leveraging a bi-directional formulation, SymmFlow enables efficient semantic reasoning while preserving the flexibility required for high-fidelity generation. Unlike prior approaches that impose rigid one-to-one mappings, it supports various conditioning strategies, including pixel- and image-level supervision. Experimental results show that SymmFlow achieves state-of-the-art performance in semantic image synthesis with only 25 inference steps, competitive segmentation accuracy despite operating in a low-resolution latent space, and promising classification results. These results show that flow-based generative models can simultaneously support generation and discrimination within a unified framework. Future work will further explore its classification potential and extend the framework to problems such as depth estimation and semantic image editing.

## Acknowledgments

This work used the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-13640. This research was funded by the European Xecs Eureka TASTI Project.

## References

- Amit, T.; Shaharbany, T.; Nachmani, E.; and Wolf, L. 2021. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*.
- Barbu, A.; Mayo, D.; Alverio, J.; Luo, W.; Wang, C.; Gutfreund, D.; Tenenbaum, J.; and Katz, B. 2019. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Adv. in Neu. Inf. Proc. Sys.*, 32.
- Black, K.; Brown, N.; Driess, D.; Esmail, A.; Equi, M.; Finn, C.; Fusai, N.; Groom, L.; Hausman, K.; Ichter, B.; et al. 2024.  $\pi_0$ : A Vision-Language-Action Flow Model for General Robot Control. *arXiv preprint arXiv:2410.24164*.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, Zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, 446–461. Springer.
- Caesar, H.; Uijlings, J.; and Ferrari, V. 2018. Coco-stuff: Thing and stuff classes in context. In *Proc. of the IEEE Conf. on Comp. Vision and Patt. Recog.*, 1209–1218.
- Chen, R. T.; Rubanova, Y.; Bettencourt, J.; and Duvenaud, D. K. 2018. Neural ordinary differential equations. *Adv. in Neu. Inf. Proc. Sys.*, 31.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girdhar, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 1290–1299.
- Cheng, B.; Schwing, A.; and Kirillov, A. 2021. Per-pixel classification is not all you need for semantic segmentation. *Adv. in Neu. Inf. Proc. Sys.*, 34: 17864–17875.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, 248–255. Ieee.
- Deng, L. 2012. The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]. *IEEE Signal Processing Magazine*, 29(6): 141–142.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Adv. in Neu. Inf. Proc. Sys.*, 34: 8780–8794.
- Ding, H.; Liu, C.; He, S.; Jiang, X.; and Loy, C. C. 2023. MeViS: A large-scale benchmark for video segmentation with motion expressions. In *Proc. of the IEEE/CVF Int. Conf. on Comp. Vision*, 2694–2703.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*.
- Gao, H.-a.; Gao, M.; Li, J.; Li, W.; Zhi, R.; Tang, H.; and Zhao, H. 2024. SCP-Diff: Spatial-Categorical Joint Prior for Diffusion Based Semantic Image Synthesis. In *European Conference on Computer Vision*, 37–54. Springer.
- Grathwohl, W.; Chen, R. T.; Bettencourt, J.; Sutskever, I.; and Duvenaud, D. 2018. Ffjord: Free-form continuous dynamics for scalable reversible generative models. *arXiv preprint arXiv:1810.01367*.
- Gu, Z.; Chen, H.; and Xu, Z. 2024. Diffusioninst: Diffusion model for instance segmentation. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2730–2734. IEEE.
- Gui, M.; Schusterbauer, J.; Prestel, U.; Ma, P.; Kotovenko, D.; Grebenkova, O.; Baumann, S. A.; Hu, V. T.; and Ommer, B. 2024. Depthfm: Fast monocular depth estimation with flow matching. *arXiv preprint arXiv:2403.13788*.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 16000–16009.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. of the IEEE Conf. on Comp. Vision and Patt. Recog.*, 770–778.
- Hinton, G. E. 2007. To recognize shapes, first learn to generate images. *Progress in brain research*, 165: 535–547.
- Hinton, G. E.; Osindero, S.; and Teh, Y.-W. 2006. A fast learning algorithm for deep belief nets. *Neural computation*, 18(7): 1527–1554.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Adv. in Neu. Inf. Proc. Sys.*, 33: 6840–6851.
- Huang, Y.; Gornet, J.; Dai, S.; Yu, Z.; Nguyen, T.; Tsao, D.; and Anandkumar, A. 2020. Neural networks with recurrent generative feedback. *Adv. in Neu. Inf. Proc. Sys.*, 33: 535–545.
- Huguet, G.; Vuckovic, J.; Fatras, K.; Thibodeau-Laufer, E.; Lemos, P.; Islam, R.; Liu, C.-H.; Rector-Brooks, J.; Akhound-Sadegh, T.; Bronstein, M.; et al. 2024. Sequence-augmented se (3)-flow matching for conditional protein backbone generation. *arXiv preprint arXiv:2405.20313*.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proc. of the IEEE Conf. on Comp. Vision and Patt. Recog.*, 1125–1134.
- Karras, T.; Aittala, M.; Aila, T.; and Laine, S. 2022. Elucidating the design space of diffusion-based generative models. *Adv. in Neu. Inf. Proc. Sys.*, 35: 26565–26577.
- Karras, T.; Aittala, M.; Laine, S.; Härkönen, E.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2021. Alias-free generative adversarial networks. *Adv. in Neu. Inf. Proc. Sys.*, 34: 852–863.

- Ko, J.; Kong, I.; Park, D.; and Kim, H. J. 2024. Stochastic conditional diffusion models for robust semantic image synthesis. *arXiv preprint arXiv:2402.16506*.
- Krizhevsky, A.; Nair, V.; and Hinton, G. 2010. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/kriz/cifar.html>, 5(4): 1.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Adv. in Neu. Inf. Proc. Sys.*, 25.
- Lee, C.-H.; Liu, Z.; Wu, L.; and Luo, P. 2020. Maskgan: Towards diverse and interactive facial image manipulation. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 5549–5558.
- Li, A. C.; Prabhudesai, M.; Duggal, S.; Brown, E.; and Pathak, D. 2023a. Your diffusion model is secretly a zero-shot classifier. In *Proc. of the IEEE/CVF Int. Conf. on Comp. Vision*, 2206–2217.
- Li, B.; Xue, K.; Liu, B.; and Lai, Y.-K. 2023b. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 1952–1961.
- Li, K.; Zhang, T.; and Malik, J. 2019. Diverse image synthesis from semantic layouts via conditional imle. In *Proc. of the IEEE/CVF Int. Conf. on Comp. Vision*, 4220–4229.
- Lipman, Y.; Chen, R. T.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2022. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*.
- Lipman, Y.; Havasi, M.; Holderrith, P.; Shaul, N.; Le, M.; Karrer, B.; Chen, R. T.; Lopez-Paz, D.; Ben-Hamu, H.; and Gat, I. 2024. Flow matching guide and code. *arXiv preprint arXiv:2412.06264*.
- Liu, X.; Gong, C.; and Liu, Q. 2022. Flow straight and fast: Learning to generate and transfer data with rectified flow. *arXiv preprint arXiv:2209.03003*.
- Liu, X.; Yin, G.; Shao, J.; Wang, X.; et al. 2019. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. *Adv. in Neu. Inf. Proc. Sys.*, 32.
- Lv, Z.; Li, X.; Niu, Z.; Cao, B.; and Zuo, W. 2022. Semantic-shape adaptive feature modulation for semantic image synthesis. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 11214–11223.
- Narayan, K.; Vs, V.; and Patel, V. M. 2025. Segface: Face segmentation of long-tail classes. In *Proc. of the AAAI Conference on Artificial Intelligence*, volume 39, 6182–6190.
- Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic image synthesis with spatially-adaptive normalization. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 2337–2346.
- Radford, A.; Metz, L.; and Chintala, S. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*.
- Ranzato, M.; Susskind, J.; Mnih, V.; and Hinton, G. 2011. On deep generative models with applications to recognition. In *CVPR 2011*, 2857–2864. IEEE.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 10684–10695.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention—MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, 234–241. Springer.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Vyas, A.; Shi, B.; Le, M.; Tjandra, A.; Wu, Y.-C.; Guo, B.; Zhang, J.; Zhang, X.; Adkins, R.; Ngan, W.; et al. 2023. Audiobox: Unified audio generation with natural language prompts. *arXiv preprint arXiv:2312.15821*.
- Wang, C.; Li, X.; Ding, H.; Qi, L.; Zhang, J.; Tong, Y.; Loy, C. C.; and Yan, S. 2024a. Explore in-context segmentation via latent diffusion models. *arXiv preprint arXiv:2403.09616*.
- Wang, C.; Li, X.; Qi, L.; Ding, H.; Tong, Y.; and Yang, M.-H. 2024b. Semflow: Binding semantic segmentation and image synthesis via rectified flow. *Adv. in Neu. Inf. Proc. Sys.*, 37: 138981–139001.
- Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proc. of the IEEE Conf. on Comp. Vision and Patt. Recog.*, 8798–8807.
- Wang, W.; Bao, J.; Zhou, W.; Chen, D.; Chen, D.; Yuan, L.; and Li, H. 2022. Semantic image synthesis via diffusion models. *arXiv preprint arXiv:2207.00050*.
- Wang, Y.; Qi, L.; Chen, Y.-C.; Zhang, X.; and Jia, J. 2021. Image synthesis via semantic composition. In *Proc. of the IEEE/CVF Int. Conf. on Comp. Vision*, 13749–13758.
- Wu, J.; Ji, W.; Fu, H.; Xu, M.; Jin, Y.; and Xu, Y. 2024. Medsegdiff-v2: Diffusion-based medical image segmentation with transformer. In *Proc. of the AAAI conference on artificial intelligence*, volume 38, 6030–6038.
- Xie, E.; Wang, W.; Yu, Z.; Anandkumar, A.; Alvarez, J. M.; and Luo, P. 2021. SegFormer: Simple and efficient design for semantic segmentation with transformers. *Adv. in Neu. Inf. Proc. Sys.*, 34: 12077–12090.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proc. of the IEEE/CVF Int. Conf. on Comp. Vision*, 3836–3847.
- Zheng, Q.; Deng, J.; Zhu, Z.; Li, Y.; and Zafeiriou, S. 2022. Decoupled multi-task learning with cyclical self-regulation for face parsing. In *Proc. of the IEEE/CVF Conf. on Comp. Vision and Patt. Recog.*, 4156–4165.
- Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. of the IEEE international conference on computer vision*, 2223–2232.