

Stochastic Segmentation with Conditional Categorical Diffusion Models

Lukas Zbinden* Lars Doorenbos* Theodoros Pissas
Adrian Thomas Huber Raphael Sznitman Pablo Márquez-Neila

University of Bern, Bern, Switzerland

{lukas.zbinden, lars.doorenbos, theodoros.pissas, raphael.sznitman, pablo.marquez}@unibe.ch

Abstract

Semantic segmentation has made significant progress in recent years thanks to deep neural networks, but the common objective of generating a single segmentation output that accurately matches the image’s content may not be suitable for safety-critical domains such as medical diagnostics and autonomous driving. Instead, multiple possible correct segmentation maps may be required to reflect the true distribution of annotation maps. In this context, stochastic semantic segmentation methods must learn to predict conditional distributions of labels given the image, but this is challenging due to the typically multimodal distributions, high-dimensional output spaces, and limited annotation data. To address these challenges, we propose a conditional categorical diffusion model (CCDM) for semantic segmentation based on Denoising Diffusion Probabilistic Models. Our model is conditioned to the input image, enabling it to generate multiple segmentation label maps that account for the aleatoric uncertainty arising from divergent ground truth annotations. Our experimental results show that CCDM achieves state-of-the-art performance on LIDC, a stochastic semantic segmentation dataset, and outperforms established baselines on the classical segmentation dataset Cityscapes.

1. Introduction

Semantic segmentation has significantly progressed in recent years due to powerful deep neural networks. For most methods, the key objective is to generate a single segmentation output that accurately matches the image’s content. However, this may not be suitable for safety-critical domains such as medical diagnostics and autonomous driving, as images in these applications often suffer from inherent ambiguity or annotations that have differences in opinion. In these cases, generating a single coherent segmentation may be hopeless to fully describe the set of correct

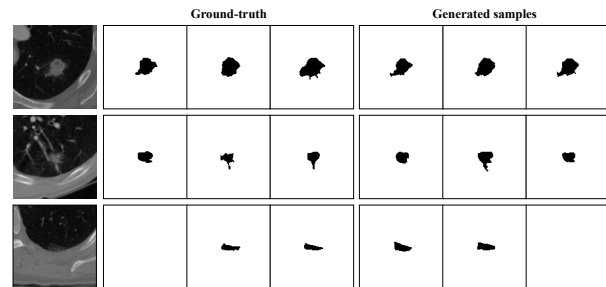


Figure 1: Examples from the LIDC dataset, where expert radiologists were asked to annotate lung nodules. Despite their expertise, they disagree significantly on many cases. Standard segmentation networks fail to capture these variations, thereby giving a false sense of confidence in model predictions. Our approach learns the distribution of possible labels, allowing us to generate realistic and diverse segmentations.

labeling.

Instead, multiple possible correct segmentation maps may be required to reflect the true distribution of annotations. For instance, Fig. 1 illustrates the task of lung nodule segmentation from CT scans where expert annotators provide multiple valid segmentation maps. In this context, stochastic semantic segmentation methods must learn to predict conditional distributions of labels given the image. Doing so is challenging, however, as the distribution is typically multimodal, the output space is high-dimensional, and annotation data is limited.

Denoising Diffusion Probabilistic Models (DDPMs) appear well-suited to overcome these challenges. DDPMs have recently drawn strong interest in computer vision as a framework for learning complex distributions in high-dimensional spaces. After achieving state-of-the-art performance on image synthesis [13], they have been successfully extended to solve tasks such as text-to-image generation [41], counterfactual explanation generation [24], inpainting [34], but also image classification [56] and seman-

*Equal contribution

tic segmentation [1, 3, 48] amongst others.

While DDPMs were originally formulated as probabilistic models able to learn high-dimensional data distributions of discrete and ordered variables (*e.g.*, RGB pixel values), re-formulations and modifications that allow for categorical variables (*e.g.*, labels) [21] are one of the key reasons why DDPMs are being explored in a broad range of computer vision tasks [12]. Specifically, the ability to model the spatial distribution of categorical variables is well suited for numerous computer vision tasks, including semantic segmentation [6, 8, 10, 14, 16, 17, 27, 31, 33, 54, 55]. Yet until now, segmentation methods using DDPMs have relied on the original discrete and ordered formulation and different heuristics to yield categorical outputs [1, 3, 48]. Consequently, the potential advantages of adopting diffusion models of categorical variables for stochastic image segmentation are still unknown.

In light of the above, we propose a *conditional categorical diffusion model* (CCDM) for semantic segmentation based on DDPMs, which models both the observed and the latent variables as categorical distributions. This enables the model to explicitly generate labels maps of discrete, unordered variables, thereby circumventing the need for switching between continuous and discrete domains, as in previous methods. The model is conditioned to the input image, making it possible to generate multiple segmentation label maps that account for the aleatoric uncertainty arising from image ambiguity. We show experimentally that our approach achieves state-of-the-art performance on LIDC, a stochastic semantic segmentation dataset, according to several performance measures. Moreover, when applied to the classical segmentation dataset Cityscapes, our method provides competitive results, outperforming established baselines.

In summary, our main contributions are the following:

- We propose a conditional categorical diffusion model capable of learning the label distribution given an input image that can be used to produce diverse segmentation samples that capture aleatoric uncertainty.
- For the task of learning a multi-rater semantic segmentation label distribution, our method achieves state-of-the-art performance on LIDC, being the first diffusion-based approach proposed for this task.
- We report competitive performance on a challenging semantic segmentation task, Cityscapes, outperforming several established baselines using a lightweight model that also leverages an off-the-shelf pre-trained feature extractor.

2. Related work

Stochastic segmentation: Methods for stochastic se-

matic segmentation aim at capturing the aleatoric uncertainty and inherent unpredictability of the labels used for segmentation. Different frameworks have been proposed to yield segmentations according to the underlying label distribution.

Initial works aimed at equipping a standard U-Net [40] with a probabilistic element to generate multiple predictions for the same image, typically accomplished by adding a conditional variational autoencoder (cVAE) [45], where the low-dimensional latent space of the cVAE encodes the possible segmentation variants. In [28], samples from this latent space are upsampled and concatenated at the last layer of the U-Net. Multiple methods extend this set-up to a hierarchical version [4, 29, 53]. Other works use normalizing flows to allow for a more expressive distribution than the Gaussian distribution in the cVAE [43, 46], switch to a discrete latent space [37], or add variational dropout and use the inter-grader variability directly as a training target [23].

Several other methods do not rely on the probabilistic U-Net. Monteiro *et al.* [35] propose a network that uses a low-rank multivariate normal distribution to model the logit distribution. Kassapis *et al.* [25] leverage adversarial training to learn possible label maps based on the logits of a trained segmentation network. Zhang *et al.* [52] employ an autoregressive PixelCNN to model the conditional distribution between pixels. Finally, Gao *et al.* [15] use a mixture of stochastic experts, where each expert network estimates a mode of the uncertainty, and a gating network predicts the probabilities that an input image is segmented by one of the experts. Our method is the first to explore the use of categorical diffusion models for stochastic segmentation.

Diffusion models: Generative diffusion models [44] have drawn much attention following their popularization by [19]. Since then, diffusion models have been successfully applied to various domains, such as image generation, restoration, and super-resolution [12].

More central to the work presented here, a few methods have attempted to apply diffusion models to semantic segmentation. Baranchuk *et al.* [3] first train diffusion models to generate images, then use multilayer perceptrons (MLP) on its features to predict the class label. Other works focus on binary segmentation with conditional diffusion models [1, 48]. These methods generate single-channel continuous samples conditioned on the input image and obtain binary segmentation masks by thresholding the result. Directly applying continuous diffusion is also done in [49, 50]. Chen *et al.* [9] generate discrete data with continuous diffusion models by encoding categorical data into bits and modeling these bits as real numbers.

Hoogeboom *et al.* [21] propose multinomial diffusion, a variation of diffusion models designed for categorical data. Subsequently, multinomial diffusion has been applied to discrete use cases, such as for tabular data [30], the latent

space of vector-quantized variational auto-encoders [11, 22] or text [21]. They can also generate segmentation maps in the unconditional setting at a very small resolution (32×64) [21]. Instead, we focus on the unexplored conditional case and demonstrate results at significantly higher resolutions (up to 256×512).

3. Method

We now introduce our approach by first framing the problem setting and defining the necessary notation. We then describe categorical diffusion models and the conditioning procedure to produce stochastic semantic segmentation via diffusion.

3.1. Background and notation

A denoising diffusion probabilistic model (DDPM) is a latent variable model $p_\theta(\mathbf{x}_0) = \int p_\theta(\mathbf{x}_{0:T}) d\mathbf{x}_{1:T}$ describing the distribution of an observable variable $\mathbf{x}_0 \in \mathbb{R}^D$ using a collection of T latent variables $\{\mathbf{x}_t\}_{t=1}^T$ with the same dimensionality as \mathbf{x}_0 . The joint distribution is modeled as a Markov chain $p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$, which is commonly known as the *reverse process*. The initial $p(\mathbf{x}_T)$ is set to a known, tractable distribution such as the Gaussian distribution, while the transition distribution p_θ , parameterized by θ , is the trainable component of the model. Training a DDPM aims to approximate $p_\theta(\mathbf{x}_0)$ to an empirical distribution $q(\mathbf{x}_0)$ defined by a collection of samples (e.g., images from the real world). To that end, training minimizes the cross-entropy between both distributions,

$$\min_{\theta} \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0)} [-\log p_\theta(\mathbf{x}_0)], \quad (1)$$

which is intractable as it requires marginalizing over the latent variables. Instead, a tractable distribution $q(\mathbf{x}_{1:T} | \mathbf{x}_0)$ is introduced and used as an approximation to the intractable true posterior $p(\mathbf{x}_{1:T} | \mathbf{x}_0)$ to define the evidence lower bound (ELBO),

$$\log p_\theta(\mathbf{x}_0) \geq \mathbb{E}_{\mathbf{x}_{1:T} \sim q(\mathbf{x}_{1:T} | \mathbf{x}_0)} \left[\log \frac{p_\theta(\mathbf{x}_{0:T})}{q(\mathbf{x}_{1:T} | \mathbf{x}_0)} \right], \quad (2)$$

where the expectation is approximated by Monte Carlo sampling. The lower bound is tight when the approximate posterior q equals the real posterior. Maximizing the ELBO over samples from $q(\mathbf{x}_0)$ minimizes the cross-entropy loss of Eq. (1).

The key difference between DDPMs and other latent variable models is that the approximate posterior $q(\mathbf{x}_{1:T} | \mathbf{x}_0)$ is fixed and not learnable. DDPMs model this distribution as a Markov chain $q(\mathbf{x}_{1:T} | \mathbf{x}_0) = \prod_{t=1}^T q(\mathbf{x}_t | \mathbf{x}_{t-1})$, known as the *forward process*. The transition distribution $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ is chosen to be a tractable distribution that allows efficient sampling from $q(\mathbf{x}_t | \mathbf{x}_0)$ for

any t . The only constraint in the design of a DDPM is that $q(\mathbf{x}_T | \mathbf{x}_0) \approx p(\mathbf{x}_T)$.

The original DDPM [19] modeled the transition distributions of the forward and the reverse processes as Gaussian with diagonal covariance matrices, and $p(\mathbf{x}_T)$ as a standard multivariate normal. However, these assumptions are inadequate when the elements of \mathbf{x}_0 belong to discrete, unordered sets, as in the task of image segmentation.

3.2. Categorical diffusion model

We now consider the denoising diffusion formulation to learn complex distributions of discrete image labelings. The observable variable $\mathbf{x}_0 \in \mathcal{L}^D$ is categorical, where D is the number of pixels of the image and $\mathcal{L} = \{1, \dots, L\}$ is the set of discrete labels that can be assigned to each pixel. Following [21], we consider that all latent variables in $\mathbf{x}_{1:T}$ are also categorical and that the transition distributions for the forward and reverse processes are modeled as categorical distributions. For the forward process, the transition distribution acts element-wise over the previous state \mathbf{x}_{t-1} to produce the parameters of the distribution for \mathbf{x}_t as,

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \prod_{d=1}^D q(\mathbf{x}_t[d] | \mathbf{x}_{t-1}[d]), \quad (3)$$

where $\mathbf{x}_t[d]$ indicates the label at time t and pixel d . In the following discussion, we will use $x_t \in \mathcal{L}$ to refer to the label of a single pixel d , and we will drop the index d for clarity. The pixel-wise transition distribution $q(x_t | x_{t-1})$ gives the element-wise probability of the next label given the previous label as,

$$q(x_t | x_{t-1}) = \mathcal{C} \left(x_t; \frac{\beta_t}{L} \mathbf{1} + (1 - \beta_t) \mathbf{e}_{x_{t-1}} \right), \quad (4)$$

where $\mathbf{1} = (1, \dots, 1)^T$, \mathbf{e}_ℓ is the one-hot encoding vector with 1 in position ℓ and 0 elsewhere, and the hyperparameter $\alpha_t = 1 - \beta_t \in (0, 1)$ indicates the probability of keeping the label unchanged. $\mathcal{C}(x; \mathbf{p})$ denotes the categorical distribution with parameter vector $\mathbf{p} \in [0, 1]^L$. From the properties of categorical distributions, $\mathcal{C}(x | \mathbf{p}) = \mathbf{p}[x]$ and $\sum_x \mathbf{p}[x] = 1$.

The transition distribution of the forward process can be composed as,

$$q(x_t | x_0) = \mathcal{C} \left(x_t; \frac{1 - \bar{\alpha}_t}{L} \mathbf{1} + \bar{\alpha}_t \mathbf{e}_{x_0} \right) \quad (5)$$

with $\bar{\alpha}_t = \prod_{\tau=1}^t \alpha_\tau$, which enables efficient sampling of elements from the Markov chain at any location t . Finally, the posterior of the transition distribution can be computed with the previous formulas by applying Bayes rule,

$$q(x_{t-1} | x_t, x_0) = \mathcal{C}(x_{t-1}; \boldsymbol{\pi}(x_t, x_0)), \quad (6)$$

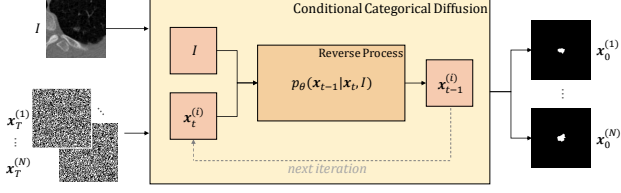


Figure 2: Illustration of the reverse process of our method. The conditional categorical diffusion model (CCDM) receives as input an image I and a categorical label map $\mathbf{x}_T^{(i)}$ sampled from the categorical uniform noise. The reverse process of the CCDM generates a label map $\mathbf{x}_0^{(i)}$, which is a sample from the learned distribution $p(\mathbf{x}_0 | I)$. When repeated for N samples, we obtain an empirical approximation to the multimodal label distribution for the image I , learned from the annotations of multiple expert raters.

with,

$$\pi(x_t, x_0) = \frac{1}{\tilde{\pi}} \left(\frac{\beta_t}{L} \mathbf{1} + \alpha_t \mathbf{e}_{x_t} \right) \odot \left(\frac{1 - \bar{\alpha}_{t-1}}{L} \mathbf{1} + \bar{\alpha}_{t-1} \mathbf{e}_{x_0} \right) \quad (7)$$

and $\tilde{\pi} = \frac{1 - \bar{\alpha}_t}{L} + \bar{\alpha}_t \cdot \delta_{x_t}^{x_0}$, where δ is the Kronecker delta.

The transition distribution of the reverse process is also an element-wise categorical distribution,

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \prod_{d=1}^D \mathcal{C}(x_{t-1}; \hat{\mathbf{p}}_{t-1}), \quad (8)$$

where $x_{t-1} = \mathbf{x}_{t-1}[d]$ and $\hat{\mathbf{p}}_{t-1}$ are the label and the estimated parameter vector, respectively, at pixel d . Unlike the forward process, the parameter vector for the pixel d is not computed considering only the element d of \mathbf{x}_t . Instead, it is modeled as a function $f : \mathcal{L}^D \rightarrow [0, 1]^{D \times L}$ that incorporates context by considering the entire label map \mathbf{x}_t to produce a collection of D probability distributions for \mathbf{x}_{t-1} , which we refer to as $\hat{\mathbf{P}}_{t-1} \in [0, 1]^{D \times L}$ with $\hat{\mathbf{p}}_{t-1} = \hat{\mathbf{P}}_{t-1}[d]$.

While it is possible to use a neural network to estimate $\hat{\mathbf{P}}_{t-1}$, Ho *et al.* [19] suggested that a consistent output space for the network led to enhanced performance. Following this idea, we train a network f_θ , parameterized by θ , to compute $\hat{\mathbf{P}}_0 = f_\theta(\mathbf{x}_t, t) \in [0, 1]^{D \times L}$ by receiving a label map \mathbf{x}_t and the step t . We then transform the parameter vector for each pixel, $\hat{\mathbf{p}}_0 = \hat{\mathbf{P}}_0[d]$ to the parameter vector $\hat{\mathbf{p}}_{t-1}$ for the same pixel of \mathbf{x}_{t-1} as,

$$\mathcal{C}(x_{t-1}; \hat{\mathbf{p}}_{t-1}) = \quad (9)$$

$$= \sum_{x_0} q(x_{t-1} | x_t, x_0) \cdot \mathcal{C}(x_0; \hat{\mathbf{p}}_0) \quad (10)$$

$$= \sum_{x_0} \mathcal{C}(x_{t-1}; \pi(x_t, x_0)) \cdot \mathcal{C}(x_0; \hat{\mathbf{p}}_0), \quad (11)$$

from which,

$$\hat{\mathbf{p}}_{t-1} = \sum_{x_0 \in \mathcal{L}} \pi(x_t, x_0) \cdot \hat{\mathbf{p}}_0[x_0], \quad (12)$$

where we have omitted the pixel indices d for clarity. This transformation is not necessary when $t = 1$, as then $\hat{\mathbf{p}}_{t-1} = \hat{\mathbf{p}}_0$ computed by f_θ . It is also possible to perform this computation in parallel for every pixel to efficiently obtain $\hat{\mathbf{P}}_{t-1}$. Note that the result of Eq. (12) differs from the parameter vector computed in [21], where the ill-defined expression $\hat{\mathbf{p}}_{t-1} = \pi(x_t, \hat{\mathbf{x}}_0)$ is employed.

3.3. Conditional categorical diffusion

In stochastic segmentation, the label map \mathbf{x}_0 for an image I is modeled by a distribution $q(\mathbf{x}_0 | I)$. This distribution is often too complex to be properly approximated as a product of pixel-wise categorical distributions. We use a conditional categorical diffusion model $p(\mathbf{x}_0 | I)$ (CCDM) to model the potentially complex interactions between labels and pixels.

When conditioning the categorical diffusion model on an image, the forward process remains unchanged, $q(\mathbf{x}_{1:T} | \mathbf{x}_0, I) = q(\mathbf{x}_{1:T} | \mathbf{x}_0)$, as any latent variable is conditionally independent of the image given any previous variable. On the other hand, the reverse process needs to incorporate the dependency on the image in its transition distribution, $p_\theta(\mathbf{x}_{0:T} | I) = p(\mathbf{x}_T | I) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, I)$. In practice, this dependency is enforced by an additional input to the neural network $f_\theta(\mathbf{x}_t, t, I)$.

3.4. Training

Training is performed by maximizing the ELBO of Eq. (2). Reorganizing terms and distributing expectations for variance reduction, we express the ELBO as a sum of three terms:

$$\log p_\theta(\mathbf{x}_0 | I) \geq \mathbb{E}_{\mathbf{x}_1 \sim q(\mathbf{x}_1 | \mathbf{x}_0)} [\log p_\theta(\mathbf{x}_0 | \mathbf{x}_1, I)] \quad (13)$$

$$- \sum_{t=2}^T \mathbb{E}_{\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0)} [KL(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, I))] \quad (14)$$

$$- KL(q(\mathbf{x}_T | \mathbf{x}_0) \| p(\mathbf{x}_T | I)). \quad (15)$$

The first two terms can be optimized by standard gradient ascent. We approximate the expectations with Monte Carlo sampling with a single sample. The sum over the time variable t is also approximated by a single uniform sample over $\{1, \dots, T\}$. The KL divergence of the second term is the sum of pixel-wise KL divergences,

$$KL(q \| p) = \sum_{d=1}^D KL(q(x_{t-1} | x_t, x_0) \| p_\theta(x_{t-1} | \mathbf{x}_t, I)), \quad (16)$$

Algorithm 1 Training a CCDM with T steps

Require: Training data expressed as the empirical distribution $q(\mathbf{x}_0, I) = q(\mathbf{x}_0 | I)q(I)$.

repeat

$t \sim \text{Uniform}(\{1, \dots, T\})$

$I \sim q(I)$

$\mathbf{x}_0 \sim q(\mathbf{x}_0 | I)$

$\mathbf{x}_t \sim q(\mathbf{x}_t | \mathbf{x}_0)$

$\hat{\mathbf{P}}_0 \leftarrow f_\theta(\mathbf{x}_t, I, t)$ ▷ shape $D \times L$

if $t > 1$ **then**

▷ Pixel-wise application of Eq. (12)

$\hat{\mathbf{p}}_{t-1} \leftarrow \sum_{x_0 \in \mathcal{L}} \pi(x_t, x_0) \cdot \hat{\mathbf{p}}_0[x_0]$ ▷ shape L

▷ Compute KL with Eq. (8) and (16)

$\ell \leftarrow KL(q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) || p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t, I))$

else

$\ell \leftarrow -\sum_d \log \mathcal{C}(x_0 | \hat{\mathbf{P}}_0[d])$

end if

$\theta \leftarrow \theta - \nabla_\theta \ell$ ▷ Gradient descent

until converged

where the parameter vectors of distributions q and p are computed with Eqs. (7) and (12), respectively. Alg. 1 shows the complete training procedure.

The third term of Eq. (15) does not depend on the learnable parameters θ and is ignored during training. It is optimized by the design of the categorical diffusion model. Since the forward process converges as

$$\lim_{t \rightarrow \infty} q(x_t | x_0) = \mathcal{C}\left(x; \frac{\mathbf{1}}{L}\right), \quad (17)$$

we fix $p(\mathbf{x}_T | I)$ to the element-wise uniform distribution,

$$p(x_T | I) = p(x_T) = \mathcal{C}\left(x_T; \frac{\mathbf{1}}{L}\right). \quad (18)$$

This ensures that $p(\mathbf{x}_T | I) \approx q(\mathbf{x}_T | \mathbf{x}_0)$, making the third term of the ELBO close to zero.

At inference, the CCDM samples from $p(\mathbf{x}_0 | I)$ to generate label maps for a given image I , which is achieved by traversing the Markov chain of the reverse process as outlined in Alg. 2 and illustrated in Fig. 2. To minimize the noise of the generated label maps, the CCDM selects the label with maximum probability instead of sampling from $\mathcal{C}(x_0 | \hat{\mathbf{p}}_0)$ in the final step.

3.5. Architecture of f_θ

As described above, the neural network f_θ receives a label map \mathbf{x}_t , a time step t , and an image I to estimate the probability parameters for \mathbf{x}_0 . Its base design is a U-Net-like architecture [13] with self-attention modules at the three innermost layers of the encoder and the decoder [13]. The network processes the input label map represented as

Algorithm 2 Inference from a CCDM with T steps

Require: Input image I , f_θ a network trained with Alg. 1

$\mathbf{x}_T \sim \mathcal{C}^D\left(x_T; \frac{\mathbf{1}}{L}\right)$

$\mathbf{x}_{\text{prev}} \leftarrow \mathbf{x}_T$ ▷ Stores interm. and final prediction

for $t = T, \dots, 1$ **do**

$\hat{\mathbf{P}}_0 \leftarrow f_\theta(\mathbf{x}_{\text{prev}}, I, t)$

if $t > 1$ **then**

▷ Pixel-wise application of Eq. (12)

$\hat{\mathbf{p}}_{t-1} \leftarrow \sum_{x_0 \in \mathcal{L}} \pi(x_t, x_0) \cdot \hat{\mathbf{p}}_0[x_0]$

$\mathbf{x}_{\text{prev}} \sim \prod_d \mathcal{C}(x_{t-1} | \hat{\mathbf{p}}_{t-1})$

else

▷ Final prediction

$\mathbf{x}_{\text{prev}} \leftarrow \arg \max_{x_0 \in \mathcal{L}} \hat{\mathbf{P}}_0[:, x_0]$

end if

end for

a binary tensor with L channels encoding the label of each pixel as a one-hot vector. Parameters of the network are shared for all values of t . The step variable t is encoded with the standard transformer sinusoidal position embedding [19] and concatenated as additional channels to the input tensor and to the feature maps of intermediate layers. Similarly, information from the input image I is presented to the network as raw pixel values concatenated to the input tensor as additional channels. In some experiments we used a pre-trained transformer architecture Dino-ViT [5] to extract informative visual features from the image I . In those cases, the extracted features were concatenated to the feature map of the third level of the U-Net encoder, which corresponds to a spatial shape equal to $\frac{1}{8}$ the shape of the input image.

4. Experiments

In all our experiments, we set $T = 250$ and the collection of β_t are set following the cosine schedule proposed in [36]. We evaluate our method on two tasks described below.

4.1. Segmentation with multiple annotations

Dataset The Lung Image Database Consortium (LIDC) [2] binary segmentation dataset consists of 1’018 three dimensional chest CT scans of patients with lung cancer. Lung nodules of each volume are annotated by four expert raters from a pool of 12, yielding large differences in annotations in some cases. We extract nodule-centered slices from the CT volumes and treat each slice as an independent image.

While LIDC is the standard benchmark of stochastic segmentation methods to date (e.g. [4, 15, 23, 25, 28, 29, 35, 43, 53, 54]), experimental configurations (pre-processing, training/validation/test splits, metrics) vastly differ across the literature. We conduct our experiments on the two most prominent LIDC splits and report results on both separately.

Method	LIDCv1						LIDCv2			
	GED ₁₆	GED ₃₂	GED ₅₀	GED ₁₀₀	HM-IoU ₁₆	HM-IoU ₃₂	GED ₁₆	GED ₅₀	GED ₁₀₀	HM-IoU ₁₆
Prob. Unet [28]	0.310 _{±0.01} ⁻	0.303 _{±0.01} ⁺	-	0.252 _{±0.004} [†]	0.552 _{±0.00} ⁻	0.548 _{±0.00} ⁺	0.320 _{±0.030} [‡]	-	0.252 _± [‡]	0.500 _{±0.030} [‡]
HProb. Unet [29]	0.270 _{±0.01} ⁻	-	-	-	0.530 _{±0.01} ⁻	-	0.270 _{±0.010} [‡]	-	-	0.530 _{±0.01}
PhiSeg [4]	0.262 _{±0.00} ⁻	0.247 _{±0.00} ⁺	-	0.224 _{±0.004} [†]	0.586 _{±0.00} ⁻	0.595 _{±0.00} ⁺	-	-	-	-
SSN [35]	0.259 _{±0.00} ⁻	0.243 _{±0.01} ⁺	-	0.225 _{±0.002}	0.558 _{±0.00} ⁻	0.555 _{±0.01} ⁺	-	-	-	-
cFlow [43]	-	0.225 _{±0.01} ⁺	-	-	-	0.584 _{±0.00} ⁺	-	-	-	-
CAR [25]	-	-	-	0.228 _{±0.009}	-	-	0.264 _{±0.002}	<u>0.248</u> _{±0.004}	<u>0.243</u> _{±0.004}	<u>0.592</u> _{±0.005}
JProb. Unet [53]	-	0.206 _{±0.00}	-	-	-	0.647 _{±0.01}	0.262 _{±0.00}	-	-	0.585 _{±0.00}
PixelSeg [52]	0.243 _{±0.01}	-	-	-	0.614 _{±0.00}	-	<u>0.260</u> _{±0.00}	-	-	0.587 _{±0.01}
MoSE [15]	0.218 _{±0.003}	-	0.195 _{±0.002}	<u>0.189</u> _{±0.002}	0.624 _{±0.004}	-	-	-	-	-
AB [9]	<u>0.213</u> _{±0.001}	<u>0.196</u> _{±0.002}	<u>0.193</u> _{±0.002}	-	0.614 _{±0.001}	0.619 _{±0.001}	-	-	-	-
CIMD [38]	0.234 _{±0.005}	0.218 _{±0.005}	0.210 _{±0.005}	-	0.587 _{±0.001}	0.592 _{±0.002}	-	-	-	-
CCDM (<i>ours</i>)	0.212 _{±0.002}	0.194 _{±0.001}	0.187 _{±0.002}	0.183 _{±0.002}	<u>0.623</u> _{±0.002}	<u>0.631</u> _{±0.002}	0.239 _{±0.003}	0.216 _{±0.003}	0.210 _{±0.003}	0.598 _{±0.001}

Table 1: Quantitative results on LIDCv1 and LIDCv2, with the methods ordered by year. **Bold** and underlined indicate best and second best per column, respectively. Our results are over 3 seeds. For GED, lower is better; for HM-IoU, higher is better. No method, including ours, uses pre-trained weights. Results for CIMD [38] and AB [9] are ours. All other scores are taken from their original papers, except (⁺) from [53], (⁻) from [53], ([†]) from [35], ([‡]) from [25].

The first, referred to as LIDCv1, is used in [4, 15, 35, 53]. LIDCv1 comprises 15’096 slices, divided into training, validation, and testing sets with the ratio 60 : 20 : 20. The second, LIDCv2, is used in [25, 28] and contains 12’816 images with the ratio 70 : 15 : 15.

Metrics We measure the performances with the Generalised Energy Distance (GED) and the Hungarian-Matched Intersection over Union (HM-IoU) [15, 25, 29]. Both metrics measure the difference between the distributions of generated and ground-truth label maps. We denote the metrics computed with n samples using a subscript, *i.e.*, GED _{n} and HM-IoU _{n} , and we set n to common values found in the literature. Note that higher number of samples yield more precise estimates.

Baselines We compare our approach to eleven recent stochastic segmentation methods: probabilistic U-Net (Prob. Unet) [28], hierarchical probabilistic U-Net (HProb. Unet) [29], PhiSeg [4], stochastic segmentation network (SSN) [35], conditional normalizing flow (cFlow) [43], calibrated adversarial refinement (CAR) [25], joint probabilistic U-Net (JProb. Unet) [53], PixelSeg [52], mixture of stochastic experts (MoSE) [15], analog bits (AB) [9], and collectively intelligent medical diffusion (CIMD) [38].

Following standard practice, we use random horizontal and vertical flipping and random rotations of 0°, 90°, 180° and 270° for data augmentation. The resolution of the input images is 128 × 128. We trained our method with the Adam optimizer [26] until convergence of the GED metric on the validation set, a polynomial learning rate scheduling starting from 1e⁻⁴ and ending with 1e⁻⁶, and batch size of 64. We applied Polyak averaging with $\alpha = 0.99995$.

4.2. Segmentation with a single annotation

We also evaluate our method with Cityscapes, a classical multi-class segmentation dataset where each image is annotated with a single label map. It comprises 2’975 RGB images of urban scenes for training and 500 images for validation, with each image labeled using 19 possible classes.

We compare our approach to several established baselines using the validation set: DeepLabv3 [7], HRNet [47], and UPerNet [51], with both ResNet [18] and Swin [32] backbones.

Besides our standard method, which performs image conditioning by concatenating the raw pixel values as channels of the input tensor, we also included in our comparison a variant CCDM-Dino which leverages pre-trained Dino-ViT features [5] as additional conditioning concatenated to intermediate feature maps of our model’s encoder.

Experiments are conducted separately for two different image resolutions: 128 × 256 and 256 × 512. For all reported methods, we first resize the images to a fixed resolution and then apply color jittering, random flipping, and standard ImageNet intensity normalization as data augmentation. All baselines are trained for 500 epochs with a batch size of 32, with optimizers, learning rate schedules, and weight decay settings as reported in their respective publications (reported in detail in the supplementary material).

Our method was trained for 800 epochs with a batch size of 32 at 128 × 256 and of 16 at 256 × 512, using the Adam optimizer [26] with a learning rate of 1e⁻⁴ linearly decayed to 1e⁻⁶. We applied Polyak averaging with $\alpha = 0.999$.

Performance is measured with the mean intersection-over-union (mIoU). Unlike GED and HM-IoU, the metric mIoU is incompatible with multiple label maps per image. During inference, CCDM generates multiple label maps per image that are subsequently fused into a single label map for

performance assessment. We found that fusing by averaging the predicted probabilities resulted in superior performances compared to fusing by majority vote.

5. Results

5.1. LIDC

We report performances on LIDCv1 and LIDCv2 in Tab. 1 and qualitative results in Fig. 3. Due to the lack of consistent evaluation protocols, we use a total of 10 metrics, thereby covering all the baselines and allowing for direct comparisons.

Our CCDM reaches the best performance for eight out of the ten metrics, despite its relatively small size, with 9M parameters compared to, *e.g.*, the 42M parameters of MoSE. CCDM also outperforms recent continuous diffusion models for segmentation, including AB [9] (9M parameters) and CIMD [38] (24M parameters). On HM-IoU₁₆, the CCDM has a lower mean performance than MoSE by 0.001, but with only half the standard deviation. The JProb. Unet reaches a higher HM-IoU₃₂ than all other methods, despite being considerably worse for GED₃₂ than our CCDM. Furthermore, on LIDCv2, the JProb. Unet achieves only the third-best score on GED₁₆, and fourth-best on HM-IoU₁₆. This result indicates how comparing results obtained on different LIDC versions with each other can be misleading.

Fig. 3 presents qualitative results from our method. In columns (g)-(l), we see that our CCDM generates a distribution of samples that captures the annotation variability created by the four expert raters. Further, as seen in the bottom example, the CCDM also generates empty samples according to the annotations (b)-(e).

Reduced number of time steps for sampling: During inference, traversing the T steps of the reverse process makes sampling from DDPMs slow. A straightforward solution [36] involves traversing only a subset of nodes of the reverse process, $\{\mathbf{x}_{k\tau} : \tau \in \{0, \dots, T/k\}\}$, reducing the number of steps by a factor k . This technique accelerates inference at the expense of reduced performance. To illustrate the trade-offs between performance and speed, Fig. 5 presents the evolution of GED₁₆ and HM-IoU₁₆ as the number of inference steps is reduced. As expected, CCDMs perform best when the number of training and inference steps are equal, but a reasonable increase in speed without a large sacrifice in performance is possible.

5.2. Cityscapes

Experimental comparisons on Cityscapes are presented in Tab. 2, and qualitative examples are provided in Fig. 4. Experiments at 128×256 demonstrate that CCDM-Dino outperforms all other methods, even when only a single sample is used. CCDM-raw also remains competitive, being outperformed only by one baseline (UPerNet+Swin-

Method			mIoU final (best)	
Architecture	Backbone	#params	128×256	256×512
DeepLabv3 [7]	ResNet50 (✓)	39m	43.4 (44.1)	58.6 (59.2)
DeepLabv3 [7]	ResNet101 (✓)	58m	43.8 (45.5)	59.2 (59.8)
UPerNet [51]	ResNet101 (✓)	83m	45.5 (47.1)	60.7 (61.2)
HRNet [47]	w48v2 (✓)	70m	48.2 (49.5)	63.3 (64.2)
UPerNet [32]	Swin-Tiny (✓)	58m	54.2 (55.9)	<u>65.5</u> (66.0)
CCDM (ours)				
samples=1		30m	53.2	60.3
samples=5		30m	55.4	62.0
samples=10		30m	56.2	62.4
CCDM (ours) Dino ViT-S (†)				
samples=1		30m + 20M	55.5	64.0
samples=5		30m + 20M	<u>56.9</u>	65.4
samples=10		30m + 20M	57.3	65.8

Table 2: Results on Cityscapes-val for resolutions 128×256 and 256×512 . **Bold** and underlined indicate best and second best per column, respectively. (✓) and (†) indicate supervised and self-supervised pretraining of the backbone, respectively. Gray indicates pretrained, non-finetuned parameters. We report final performance for our method and baselines. For the latter we also provide best achieved performance during training (in parenthesis). For CCDM methods, the field *samples* indicates the number of generated samples for label map fusion, as explained in Sect 4.2.

CCDM Capacity		mIoU (128 × 256)		
#params	UNet Levels	samples=1	samples=5	samples=10
5.4M	4	37.8	39.7	40.6
7.5M	5	44.7	48.3	48.5
22M	4	51.6	54.0	53.6
30M	5	53.2	55.4	56.2

Table 3: Effect of increasing CCDM capacity (without feature conditioning).

Tiny), despite using only between 36% and 51% of the parameters of other models. Similarly, at 256×512 , CCDM-Dino outperforms four of the baselines with a single sample, lags behind UPerNet+Swin-Tiny only by 0.1 percent points with 5 samples, and outperforms all baselines with 10 samples. As expected, averaging across more samples improves performance for both CCDM-raw and CCDM-Dino, albeit with diminishing gains. Furthermore, the addition of Dino features boosts single-sample performance by 2.3 percent points at 128×256 , and 3.7 percent points at 256×512 , hinting the greater value of adding feature conditioning for generating segmentation at a higher resolution.

CCDM Capacity: Tab. 3(b) demonstrates the effect of increasing the capacity of CCDM. Using more U-Net feature levels, and increasing the number of parameters by doubling the number of channels per level, increases the performance regardless of the number of samples used for inference.

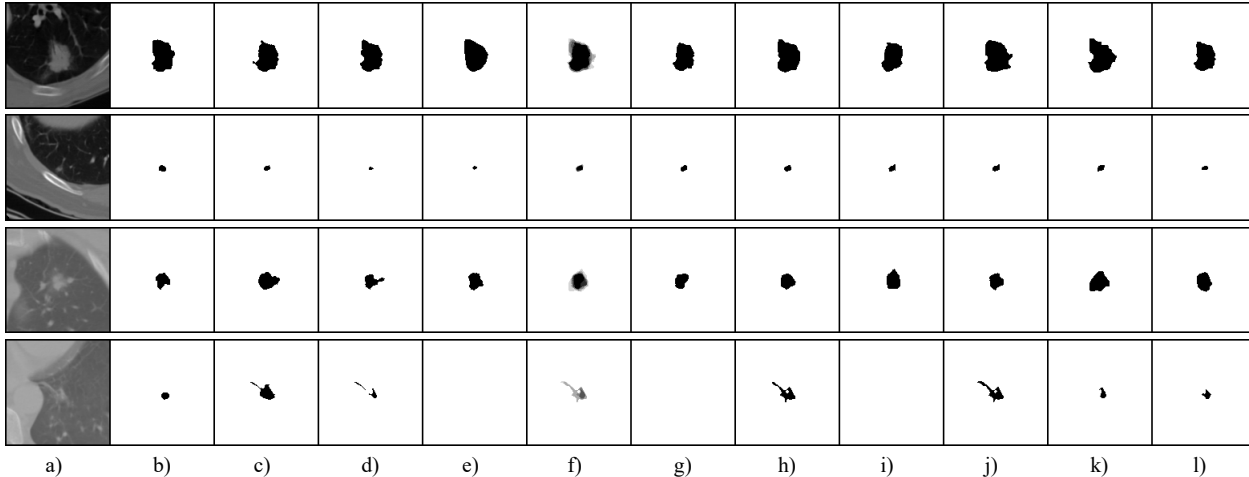


Figure 3: Qualitative results on four LIDC images with our method. (a) shows the image, (b)-(e) its four labels, (f) the mean prediction of our CCDM over six predictions, and (g)-(l) six individual predictions.

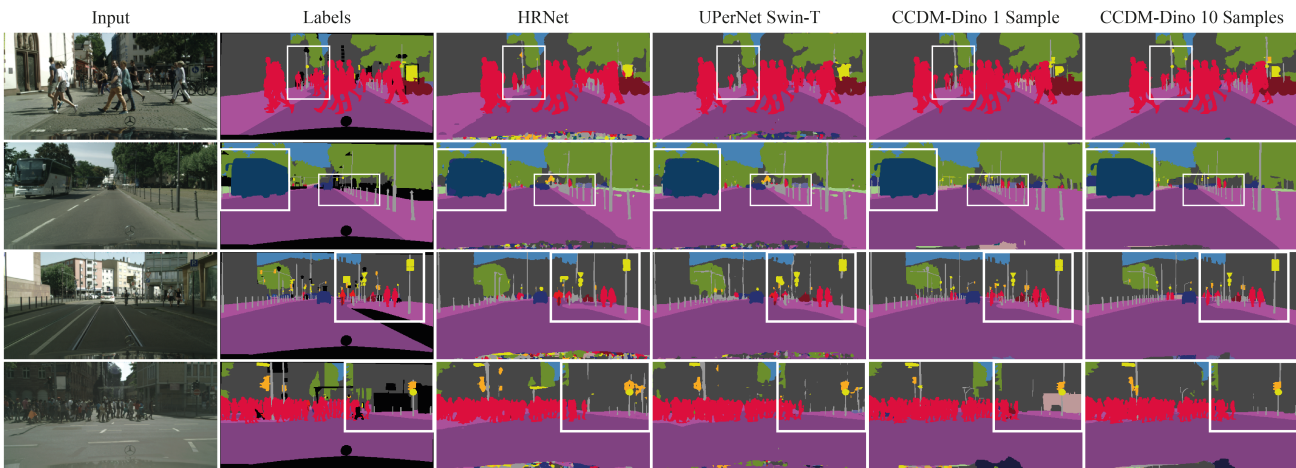


Figure 4: Qualitative comparisons on Cityscapes. All methods are trained and tested at a resolution of 256×512 . Our method produces structures with greater visual realism than other baselines. This is especially noticeable inside the marked regions.

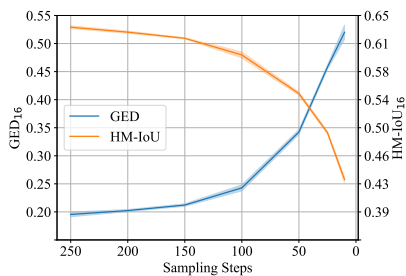


Figure 5: LIDC GED and HM-IoU versus the number of sampling steps on LIDC. Evaluated on 500 random test images using 16 samples each, over 3 seeds.

6. Conclusion

We introduced conditional categorical diffusion models (CCDMs) that are capable of effectively modeling pixel-level semantic distributions. Notably, and contrary to standard deterministic segmentation approaches, our model can produce diverse samples given an input image, thereby capturing the aleatoric uncertainty. Our method learns a multimodal label distribution of segmentations, induced by annotations from multiple expert raters, for which it achieves state-of-the-art results on a challenging medical imaging dataset, LIDC. Additionally, we demonstrate that it can achieve competitive performance on a standard multi-class semantic segmentation benchmark, Cityscapes, by outperforming several established, heavily engineered baselines

despite using significantly fewer parameters.

One limitation of our method is the requirement of several iterations for producing a sample, which is a common shortcoming of diffusion models. Accelerating sampling constitutes a crucial research direction, orthogonal to the present work. Finally, resolution scaling remains notoriously difficult for diffusion models, with successful examples relying on massive computational resources to train cascades of models that gradually increase resolution [20, 42] or operate on the latent space of existing embedding methods for continuous data (e.g. images) [39] that are not available for categorical data.

Acknowledgements

This work was partially funded by the University of Bern, Swiss National Science Foundation Grants #320030-188591, #200021-192285 and #200021-191983.

References

- [1] Tomer Amit, Eliya Nachmani, Tal Shaharby, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. 2
- [2] Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer, Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman, et al. The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011. 5
- [3] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khrukov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. 2
- [4] Christian F Baumgartner, Kerem C Tezcan, Krishna Chaitanya, Andreas M Hötker, Urs J Muehlethaler, Khoshy Schawkat, Anton S Becker, Olivio Donati, and Ender Konukoglu. Phiseg: Capturing uncertainty in medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 119–127. Springer, 2019. 2, 5, 6
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. 5, 6
- [6] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2
- [7] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 6, 7
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 2
- [9] Ting Chen, Ruixiang Zhang, and Geoffrey Hinton. Analog bits: Generating discrete data using diffusion models with self-conditioning. *arXiv preprint arXiv:2208.04202*, 2022. 2, 6, 7
- [10] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. *Advances in Neural Information Processing Systems*, 34:9355–9366, 2021. 2
- [11] Max Cohen, Guillaume Quispe, Sylvain Le Corff, Charles Ollion, and Eric Moulines. Diffusion bridges vector quantized variational autoencoders. *arXiv preprint arXiv:2202.04895*, 2022. 3
- [12] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *arXiv preprint arXiv:2209.04747*, 2022. 2
- [13] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. *CoRR*, abs/2105.05233, 2021. 1, 5
- [14] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3146–3154, 2019. 2
- [15] Zhitong Gao, Yucong Chen, Chuyu Zhang, and Xuming He. Modeling multimodal aleatoric uncertainty in segmentation with mixture of stochastic expert. *arXiv preprint arXiv:2212.07328*, 2022. 2, 5, 6
- [16] Jiaqi Gu, Hyoukjun Kwon, Dilin Wang, Wei Ye, Meng Li, Yu-Hsin Chen, Liangzhen Lai, Vikas Chandra, and David Z Pan. Multi-scale high-resolution vision transformer for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12094–12103, 2022. 2
- [17] Adam W Harley, Konstantinos G Derpanis, and Iasonas Kokkinos. Segmentation-aware convolutional networks using local attention masks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5038–5047, 2017. 2
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 6
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020. 2, 3, 4, 5
- [20] Jonathan Ho, Chitwan Saharia, William Chan, David J. Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022. 9

- [21] Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021. [2](#), [3](#), [4](#)
- [22] Minghui Hu, Yujie Wang, Tat-Jen Cham, Jianfei Yang, and Ponnuthurai N Suganthan. Global context with discrete diffusion in vector quantised modelling for image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11502–11511, 2022. [3](#)
- [23] Shi Hu, Daniel Worrall, Stefan Knecht, Bas Veeling, Henkjan Huisman, and Max Welling. Supervised uncertainty quantification for segmentation with multiple annotations. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part II 22*, pages 137–145. Springer, 2019. [2](#), [5](#)
- [24] Guillaume Jeanneret, Loïc Simon, and Frédéric Jurie. Diffusion models for counterfactual explanations. *arXiv preprint arXiv:2203.15636*, 2022. [1](#)
- [25] Elias Kassapis, Georgi Dikov, Deepak K Gupta, and Cedric Nugteren. Calibrated adversarial refinement for stochastic semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7057–7067, 2021. [2](#), [5](#), [6](#)
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference for Learning Representations*, 2015. [6](#)
- [27] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6399–6408, 2019. [2](#)
- [28] Simon Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R Ledsam, Klaus Maier-Hein, SM Eslami, Danilo Jimenez Rezende, and Olaf Ronneberger. A probabilistic u-net for segmentation of ambiguous images. *Advances in neural information processing systems*, 31, 2018. [2](#), [5](#), [6](#)
- [29] Simon AA Kohl, Bernardino Romera-Paredes, Klaus H Maier-Hein, Danilo Jimenez Rezende, SM Eslami, Pushmeet Kohli, Andrew Zisserman, and Olaf Ronneberger. A hierarchical probabilistic u-net for modeling multi-scale ambiguities. *arXiv preprint arXiv:1905.13077*, 2019. [2](#), [5](#), [6](#)
- [30] Akim Kotelnikov, Dmitry Baranchuk, Ivan Rubachev, and Artem Babenko. Tabddpm: Modelling tabular data with diffusion models. *arXiv preprint arXiv:2209.15421*, 2022. [2](#)
- [31] Liulei Li, Tianfei Zhou, Wenguan Wang, Jianwu Li, and Yi Yang. Deep hierarchical semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1246–1257, 2022. [2](#)
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021. [6](#), [7](#)
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. [2](#)
- [34] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11461–11471, 2022. [1](#)
- [35] Miguel Monteiro, Loïc Le Folgoc, Daniel Coelho de Castro, Nick Pawlowski, Bernardo Marques, Konstantinos Kamnitsas, Mark van der Wilk, and Ben Glocker. Stochastic segmentation networks: Modelling spatially correlated aleatoric uncertainty. *Advances in Neural Information Processing Systems*, 33:12756–12767, 2020. [2](#), [5](#), [6](#)
- [36] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning*, pages 8162–8171. PMLR, 2021. [5](#), [7](#)
- [37] Di Qiu and Lok Ming Lui. Modal uncertainty estimation via discrete latent representation. *arXiv preprint arXiv:2007.12858*, 2020. [2](#)
- [38] Aimon Rahman, Jeya Maria Jose Valanarasu, Ilker Hacihaliloglu, and Vishal M Patel. Ambiguous medical image segmentation using diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11536–11546, 2023. [6](#), [7](#)
- [39] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10684–10695, June 2022. [9](#)
- [40] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. [2](#)
- [41] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*, 2022. [1](#)
- [42] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. [9](#)
- [43] Raghavendra Selvan, Frederik Faye, Jon Middleton, and Akshay Pai. Uncertainty quantification in medical image segmentation with normalizing flows. In *Machine Learning in Medical Imaging: 11th International Workshop, MLMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings 11*, pages 80–90. Springer, 2020. [2](#), [5](#), [6](#)
- [44] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using

- nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015. [2](#)
- [45] Kihyuk Sohn, Honglak Lee, and Xinchun Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015. [2](#)
- [46] MM Amaan Valiuddin, Christiaan GA Viviers, Ruud JG van Sloun, Peter HN de With, and Fons van der Sommen. Improving aleatoric uncertainty quantification in multi-annotated medical image segmentation with normalizing flows. In *Uncertainty for Safe Utilization of Machine Learning in Medical Imaging, and Perinatal Imaging, Placental and Preterm Image Analysis: 3rd International Workshop, UNSURE 2021, and 6th International Workshop, PIPPI 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, October 1, 2021, Proceedings 3*, pages 75–88. Springer, 2021. [2](#)
- [47] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, Wenyu Liu, and Bin Xiao. Deep high-resolution representation learning for visual recognition. *TPAMI*, 2019. [6](#), [7](#)
- [48] Julia Wolleb, Robin Sandkühler, Florentin Bieder, Philippe Valmaggia, and Philippe C Cattin. Diffusion models for implicit image segmentation ensembles. *arXiv preprint arXiv:2112.03145*, 2021. [2](#)
- [49] Junde Wu, Huihui Fang, Yu Zhang, Yehui Yang, and Yanwu Xu. Medsegdiff: Medical image segmentation with diffusion probabilistic model. *arXiv preprint arXiv:2211.00611*, 2022. [2](#)
- [50] Junde Wu, Rao Fu, Huihui Fang, Yu Zhang, and Yanwu Xu. Medsegdiff-v2: Diffusion based medical image segmentation with transformer. *arXiv preprint arXiv:2301.11798*, 2023. [2](#)
- [51] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 418–434, 2018. [6](#), [7](#)
- [52] Wei Zhang, Xiaohong Zhang, Sheng Huang, Yuting Lu, and Kun Wang. Pixelseg: Pixel-by-pixel stochastic semantic segmentation for ambiguous medical images. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4742–4750, 2022. [2](#), [6](#)
- [53] Wei Zhang, Xiaohong Zhang, Sheng Huang, Yuting Lu, and Kun Wang. A probabilistic model for controlling diversity and accuracy of ambiguous medical image segmentation. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4751–4759, 2022. [2](#), [5](#), [6](#)
- [54] Yifan Zhang, Bo Pang, and Cewu Lu. Semantic segmentation by early region proxy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1258–1268, 2022. [2](#), [5](#)
- [55] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [2](#)
- [56] Roland S Zimmermann, Lukas Schott, Yang Song, Benjamin A Dunn, and David A Klindt. Score-based generative classifiers. *arXiv preprint arXiv:2110.00473*, 2021. [1](#)