003

004

005 006

007

008

010

011

012

013

014

015

016

018

019

020

021

022

023

024

025

026

027

028

029

030

032

033

034

035

036

037

039

040

041

042

043

044

045

048

049

051

055

056

063

066

067

069

070

071

Towards Agnostic and Holistic Universal Image Segmentation with Bit Diffusion

Anonymous Full Paper Submission 42

Abstract

This paper introduces a diffusion-based framework for universal image segmentation, making agnostic segmentation possible without depending on maskbased frameworks and instead predicting the full segmentation in a holistic manner. We present several key adaptations to diffusion models, which are important in this discrete setting. Notably, we show that a location-aware palette with our 2D gray code ordering improves performance. Adding a final tanh activation function is crucial for discrete data. On optimizing diffusion parameters, the sigmoid loss weighting consistently outperforms alternatives, regardless of the prediction type used, and we settle on x-prediction. While our current model does not yet surpass leading mask-based architectures, it narrows the performance gap and introduces unique capabilities, such as principled ambiguity modeling, that these models lack. All models were trained from scratch, and we believe that combining our proposed improvements with large-scale pretraining or promptable conditioning could lead to competitive models.

1 Introduction

In universal image segmentation, the goal is to segment images from many data modalities with a single model. Conversely, in narrow image segmentation, you specialize on a single dataset or task, such as brain tumor segmentation. In recent times, the image segmentation field has favored universal segmentation models such as Mask-RCNN [1] and the Segment Anything Model (SAM) [2, 3]. These foundation models are used as general problem solvers that can be finetuned or prompted for narrow downstream tasks.

Universal segmentation systems increasingly face two, often competing, requirements: agnostic behavior and a holistic view of the images. By agnostic we mean the ability to segment objects without relying on a fixed label set. Agnostic models focus on masks while unbound by labels, enabling the model to generalize across domains and unseen categories. By holistic we mean a model that considers the whole image when producing segmentations, including inter-mask correlations. I.e., choosing the same semantic division for separate masks. In prac-

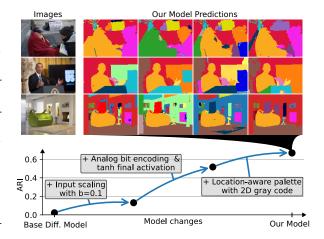


Figure 1. The modifications to a base diffusion model and their performance gains, visualized along with samples from our model.

tice, the first property enables open-world and cross-dataset use, while the second reduces segmentation inconsistencies.

We study diffusion-based segmentation as a route to achieve these goals. Diffusion models are well known for revolutionizing image generation [4], but in our setting the image is only a conditional input to the task of generating the segmentation. Additionally, using diffusion models makes ambiguity modeling possible.

Direct diffusion over discrete, high-dimensional label spaces is non-trivial. Diffusion was developed with continuous targets in mind, and it therefore faces multiple challenges when dealing with discrete data such as segmentations. Our approach combines various ideas from the diffusion research landscape. The addition of these ideas is essential to raise our model's performance. Our main contribution is to adapt the following techniques to work for diffusion segmentation (see Fig. 1):

- 1. Input scaled noise schedule [5]. We propose using an input-scaled diffusion noise schedule, in order to make the denoising problem suitably hard for discrete target spaces and improving training stability for segmentation.
- 2. Analog bit diffusion encoding [6]. We encode 2^k classes with k signed bits and train the diffusion model to predict bit-valued targets, reducing dimensionality while preserving a simple of

078

079

080

081

082

083

084

085

086

087

088

089

090

091

093

094

095

096

097

098

099

101

102

103

104

105

106

109

110

111

112

113

114

116

117

118

119

120

124

125

126

127

128

140

141

142

144

145

152

153

155

156

157

158

160

163

166

167

168

173

176

178

179

180

181

182

183

route back to class indices. We suggest adding a $\tanh(\cdot)$ activation, as it aligns the network's outputs with the discrete bit codes and yields better-calibrated probabilities.

3. Location-aware palette [7] (LAP). We adapt an LAP to reduce the downsides of the analog bit encoding when paired with our ordering that follows a 2D gray code. The LAP assigns indices by mask location, creating consistent targets in an agnostic setting, improving training.

2 Related Works

The most common flavor of universal segmentation models are mask-based (e.g. Mask R-CNN [1]). They generally work by detecting candidate regions for potential masks, and then handling each candidate separately as a binary mask prediction and/or classification problem [8–11]. Promptable class-agnostic systems such as SAM [2] demonstrate strong open-world mask extraction, but are still relying on binary foreground/background mask prediction. Masks are produced independently across the image and are therefore not holistic. An ideal universal segmentation model should be holistic, to avoid inconsistency when producing e.g. repeating objects in an image or simply to avoid overlapping masks.

Mask-based models are limited to predicting one mask at a time because they optimize for mean predictions. For full agnostic segmentations, the mean would deviate too far from any ground truth due to scene uncertainty. This issue is less severe for binary masks, where variance is low, and absent in non-agnostic models with fixed vocabularies. Traditional losses such as cross-entropy or Dice push toward single estimates even when boundaries are ill-defined or annotators disagree, often blurring details and under-representing multi-modal solutions. Probabilistic segmentation explicitly models these uncertainties, e.g., Probabilistic U-Net and its variants [12–14], and hierarchical variational approaches [15]. Bayesian [16] and ensemble-style methods estimate uncertainty but often at a significant compute cost or weaker distributional guarantees. Diffusion-based segmentation offers a generative alternative that can sample diverse, plausible masks and produce uncertainty maps by construction [17–20]. Previously mentioned generative models all operate on narrow tasks instead of universal segmentation.

Another recent diffusion-based method, pix2seq-D [21] focused on panoptic segmentation with diffusion models. They took advantage of the ambiguity modeling inherent to generative models by splitting semantic masks into instance

masks without running into combinatorial problems. Their method also made use of input scaling and analog bits, to deal with the discrete data domain.

The paper Unified Representation for Image Generation and Segmentation (UniGS) [7] is the most comparable to our approach, as it also tackles universal image segmentation with diffusion models. UniGS treats masks and images within a single latent-diffusion framework by representing entitylevel masks as RGB colormaps aligned to the image domain. They choose the RGB space because their network is a finetuned Stable Diffusion [4] model (text-to-image). Decoding masks from the predicted RGB encoding is tricky, requiring the introduction of a progressive dichotomy module. The authors also introduce a location-aware color palette that assigns consistent colors to entities based on spatial location. Relative to UniGS, our work only targets the segmentation domain and instead of utilizing a pretrained model such as Stable Diffusion, we train from scratch. Training from scratch comes with upsides and downsides, namely we are restricted to working at a small scale but we are able to study the properties of the model in an unbiased setting, and without restrictions on modeling choices.

3 Methods

3.1 Diffusion Model

We use a continuous time diffusion model [22, 23] ranging from time t = 0 (data) to t = 1 (noise). The diffusion sample x_t is given by the equation

$$\mathbf{x}_t = \alpha(t)\mathbf{x}_0 + \sigma(t)\epsilon,\tag{1}$$

where \mathbf{x}_0 is data, ϵ is i.i.d unit Gaussian noise. For a diffusion segmentation model such as ours, the data is a segmentation map. The image is a conditional input which we concatenate across the channel dimension. The model operates in pixel space, since recent research shows these models can be competitive latent diffusion alternatives [24, 25].

In order to predict \mathbf{x}_0 , the network can predict it directly (x-prediction), predict the noise (ϵ -prediction), or predict $\mathbf{v} = \alpha(t)\epsilon - \sigma(t)\mathbf{x}_0$ (v-prediction [22]). Each of these predictions parameterize the others based on Eq. (1).

We employ a convolutional neural network (CNN) with an attention mechanism to predict the mean of the conditional distribution $p(\mathbf{x}_0|\mathbf{x}_t)$, i.e. predicting the data from a noisy latent sample. Based on [23], the model can generate segmentation maps by denoising pure noise into segmentation maps over a number of timesteps. We always use equidistant timesteps from t=1 to t=0 when sampling.

The model is trained with the weighted MSE loss function [23]

$$L(\mathbf{x}) = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[w(t) \| \mathbf{x}_0 - \hat{\mathbf{x}} \|^2 \right], \qquad (2) \quad 184$$

where $\hat{\mathbf{x}} = \hat{\mathbf{x}}_{\theta}(\mathbf{x}_t, t)$ is the neural network prediction of the data, \mathbf{x}_0 . The loss weighting, w(t), can emphasize the importance of different parts of the diffusion process, and following [24, 25] we use the sigmoid loss weighting with a bias of -4.

3.2 Bit Diffusion

Analog Bit Diffusion [6] is a modification to diffusion models that enable the model to work with high-dimensional discrete data, while maintaining a low dimensional latent space. Instead of representing discrete data as e.g. one-hot vectors, we represent the 2^k classes as $n_{\rm bits}=k$ bits. We use $2^6=64$ classes corresponding to $n_{\rm bits}=6$. Negative bits have a value of -1 instead of 0, to make their distribution zero-mean and unit variance.

The diffusion process works in the bit space, and can be easily converted to the class space by thresholding the bits at 0 and converting from the binary representation.

With the bit diffusion formulation, the model should only predict values within [-1,1] with heavy emphasis on the endpoints of the interval. The $\tanh(\cdot)$ activation function is well suited for such a distribution, and we therefore apply it as a final activation (in cases where the model predicts the data directly). The non-thresholded bit activations enable conversion to a direct probability map. Let \hat{y} be the predicted bits for some pixel. The probability that the pixel has the binary sequence y is given by

$$p(y|\hat{y}) = \prod_{i=0}^{n_{\text{bits}}-1} p(y_i|\hat{y}_i) = \prod_{i=0}^{n_{\text{bits}}-1} \left(1 - \frac{|y_i - \hat{y}_i|}{2}\right).$$
(3)

The equation above makes the downside of using a bit encoding clear. It does not model correlations between bits, but instead considers each bit probability separately.

3.3 Noise Schedule and Input Scaling

The noise schedule is parameterized by $\gamma:[0,1] \to [0,1]$, a monotonically decreasing function. We use a variance preserving noise schedule, where the coefficients are given by

$$\alpha(t) = \sqrt{\gamma(t)}, \quad \sigma(t) = \sqrt{1 - \gamma(t)}.$$
 (4)

The variance preserving property enables parameterizing both set of coefficients with a single function. A common choice for the noise schedule is the cosine schedule, which is given by $\gamma(t) = \cos(t\pi/2)^2$.

Consider the upper row of latent samples in Fig. 2. As a consequence of using discrete data with high spatial correlation, it is easy to reconstruct the data for large parts of the diffusion process. If the model is able to only consider the latent sample for large parts of the diffusion process during training, then

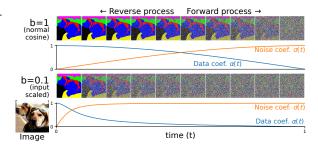


Figure 2. The cosine noise schedule with latent diffusion samples x_t for various values of t. The latent samples use 3 bits (up to 8 masks) to make them viewable as RGB images.

the resulting model will be poor since it ignores the image during inference.

To address these concerns we use input scaling [5], which can be used to make diffusion noise schedules harder. Input scaling was originally introduced to deal with large images since increasing the number of pixels lessens the effect of the noise. The idea behind input scaling is to make noise schedule harder by lowering the signal-to-noise ratio (SNR). The SNR is given by

$$SNR(t) = \frac{\alpha(t)}{\sigma(t)} = \frac{\sqrt{\gamma(t)}}{\sqrt{1 - \gamma(t)}},$$
 (5) 245

and is lowered by multiplying with some constant $b \in [0, 1]$, called the input scale. One can show that solving

$$\frac{\sqrt{\gamma_b(t)}}{\sqrt{1 - \gamma_b(t)}} = b \frac{\sqrt{\gamma(t)}}{\sqrt{1 - \gamma(t)}},\tag{6}$$

for the input scaled noise schedule, $\gamma_b(t)$, yields the expression 251

$$\gamma_b(t) = \frac{b^2 \gamma(t)}{(b^2 - 1)\gamma(t) + 1}. (7) 252$$

Thus, all equations involving the noise schedule can be reused, except by replacing the original $\gamma(t)$ with the input scaled $\gamma_b(t)$.

3.4 Location-aware Palette

The segmentation model is class-agnostic, and therefore the class numbers which we assign objects can be permuted without changing the task. A valid option is thus to assign random class numbers, but a better option is a location-aware palette (LAP)[7]. With an LAP, each mask is assigned a class number based on the mask centroid. An $L \times L$ grid is constructed across the image, with each square associated with a class number. When multiple mask centroids share a grid, they are instead given the class number of the nearest free grid square. Without an LAP, the best prediction at t=1 is a zero-image, since the data is pure noise and the expected value of random

bits is zero. When the prediction is independent of the image, there is no useful learning signal. When using an LAP, classes are biased towards the nearby LAP class indices, thus providing a learning signal for the parts of the diffusion process where the latent sample is largely noise.

The analog bit encoding has difficulty representing class distributions with multiple classes when the bits of the classes differ significantly (see supplementary material for details). By exploiting the LAP, we can increase the likelihood of adjacent class regions sharing their bit encoding digits. To this end, we arrange the bit codes in the $L \times L$ grid as a 2-dimensional gray code [26]. This ensures each 1-connectivity pair of neighbors only differ by 1 bit in the LAP. Since we use $n_{\rm bits}=6$ we have $L=\sqrt{2^6}=8$.

4 Experiments

4.1 Evaluation Setup

As a basis for our experiments we use the Entity-Seg [27] dataset, consisting of 33,227 images each fully segmented with high-quality agnostic class labels across a variety of modalities. We partition the dataset on a holdout basis with an 80-10-10 split (train-val-test) and we use a 128×128 resolution version of their dataset using the padding strategy from [2]. Our model is a 38.5m parameter attn-UNet trained for 300k iterations with a batch size of 8. The learning rate was set at 1e-4, with linear warmup for the first 1000 iterations and decreased with a cosine schedule for the last 50k iterations. We used the AdamW [28] optimizer.

We compare quantitatively using two metrics. The first is the adjusted rand index (ARI), which is based on the probability of two random pixels agreeing in the ground truth and prediction on whether they should belong to the same class or different classes. The adjusted formulation ensures the expected value for a random prediction is 0 while still keeping a perfect prediction at a score of 1. The second metric is the Intersection over Union (IoU) matched with the Hungarian algorithm. Following [27] we only compute the mean over non-empty ground truth classes after matching ground truths with predictions.

Our main model uses x-prediction and the sigmoid loss weights. The noise schedule is a cosine noise schedule with input scale parameter b=0.1. The training data class indices are chosen based on a location-aware palette (LAP) that promotes similar analog bit encodings. A final activation function of $\tanh(\cdot)$ is applied to the network. For sampling, we use 8 timesteps and a guidance weight of 1.0 unless otherwise is stated.

	No LAP		w/ LAP	
Encoding	ARI	\mathbf{IoU}	ARI	\mathbf{IoU}
Onehot	0.168	0.186	0.528	0.323
RGB	0.460	0.283	0.524	0.312
Analog Bits	0.515	0.368	0.670	0.432

Table 1. Mean performance for models trained with different encoding types.

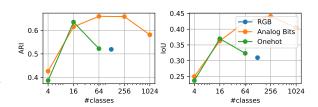


Figure 3. Performance for the three encoding types as the number of representable classes are varied.

4.2 Comparisons

We compare our model with the onehot and RGB encodings. The results (shown in Table 1) show that our model using analog bits improves upon the alternatives. The contrast is especially large when the models are trained with an LAP.

The analog bit encoding has exponential efficiency in the number of classes it can represent, which is clear when comparing how many classes the methods can represent in Fig. 3. Onehot and analog bits are similar in performance until around 16 classes when onehot falls off. We use 64 classes as a baseline for the rest of the experiments, since 96.14% of images in the dataset have less than 65 objects.

Our model was trained with an empty image in 5% of training samples, as it enables using classifier free guidance [29] during sampling to increase the conditioning strength. To optimize sampling, we vary the guidance weight (gw) and number of sampling timesteps (see Fig. 4 and Fig. 5). We see that around only 8 sampling steps is optimal and performance only degrades slightly when using more steps. Based on the ARI metric gw=1.0 is best, while IoU prefers a stronger gw=2.5. Note that gw=0.0 is the same as normal sampling with no guidance.

We increase the number of samples for each image in Fig. 6 to see the potential gains if one had an oracle to select the best prediction. More realistically this indicates the usefulness of a human in the loop or a test time augmentation (TTA) heuristic to select or aggregates samples. Using a larger guidance weight comes with a small penalty for the sample diversity as we see a smaller gain in performance.

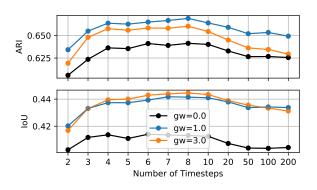


Figure 4. Mean performance on the validation set as the number of timesteps is varied for different guidance weights (gw).

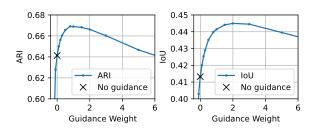


Figure 5. Mean performance on the validation set as the guidance weight is varied.

4.3 Ablations

To investigate the best pair of prediction type and loss weights, we train a range of models while varying the available options. The results are seen in Fig. 7. With all prediction types, the sigmoid loss weights perform the best. The model with ϵ -prediction is slightly better than x-prediction. However, when inspecting samples produced by the model (see Fig. 8, the ϵ -prediction often failed to remove all the noise. One might think thresholding would solve this problem, but based on qualitative inspection of samples it seems the denoising trajectory is affected, leaving small noisy patches of nonsensical labels. A much more visible symptom of the same effect is visible for the model with no tanh activation. Given the

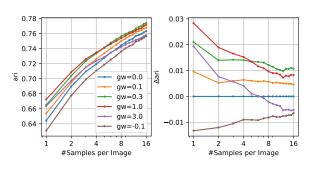


Figure 6. Mean performance when selecting the best segmentation from multiple samples. Shown in absolute ARI (left) and relative to no guidance (right).

\mathbf{LAP}	None	Different	Random	Similar
ARI	0.517	0.640	0.644	0.670
IoU	0.367	0.422	0.421	0.434

Table 2. Mean performance for models trained with different types of LAP.

tiny difference in performance, we therefore still use x-prediction.

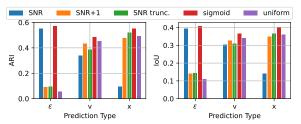


Figure 7. Mean performance for models trained with different prediction types and loss weights. These models were trained without LAP and b = 0.1.

The LAP encoding setup described in Section 3.4 is the one we call similar, since adjacent encodings are similar. Additionally, we also consider an LAP with random class indices and one which maximizes the difference of adjacent classes based on a greedy heuristic. The results in Table 2 show that in all cases, an LAP significantly increases performance. Additionally, the more similar the bit encodings of adjacent class indices, the better the performance.

To study the effect of input scaling we train a variety of models while varying b (see Fig. 9). A value of b=0.1 is close to optimal for our application. Note that b=1.0 corresponds to a model with no input scaling. The average metrics are more than doubled by just adding input scaling to the noise schedule.

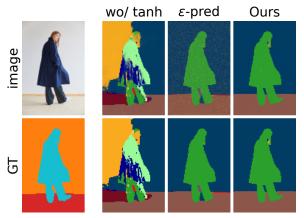


Figure 8. A qualitative example to illustrate the difference in samples produced by a model without $\tanh(\cdot)$ activation and with ϵ -prediction, compared to our model.

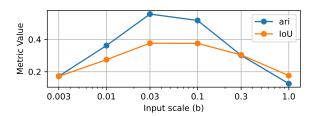


Figure 9. Mean performance for non-LAP models when varying the input scale parameter (b).

5 Discussion

Our experiments show that analog bits consistently outperforms RGB and one-hot encodings in agnostic segmentation. The relative gains are largest when class indices are assigned with a location-aware palette (LAP). We theorize that the gain in performance is an effect of an improved training process. Previously the network would learn little to nothing near t=1, just producing a zero-mean prediction, but the bias from the LAP lets it encode segmentations at any timestep. By ordering bit codes of the LAP with a 2D gray ordering, we reduced differences between neighboring bits (the similar model). This allowed the model to express soft ambiguity between adjacent masks without paying the penalty of spreading probability mass over many unrelated codes

The analog bit encoding was preferred in our networks that were trained from scratch. An interesting research question is whether the same holds for tasks similar to that of UniGS [7]. The UniGS model was designed with the RGB encoding specifically because stable diffusion operates in RGB space. It may be possible to add a head to the segmentation branch to make this conversion possible. Given UniGS already reports competitive scores in segmentation benchmarks, replacing RGB colormaps with analog bits could perhaps push the unified generator—segmenter model to the forefront.

The best results were achieved when the network used x-prediction. Across prediction types $(x, v, and \epsilon)$, the sigmoid loss weighting dominates alternatives, provided its bias is tuned. In our early experiments, we found a bias of -4 to be effective. Input scaling makes the schedule "hard enough" for discrete targets: reducing the effective SNR with $b{\approx}0.1$ more than doubles ARI over the unscaled cosine schedule. Since input scaling was introduced in order to tackle the problem of high-dimensional spatial data, one can expect it should be lowered further than b=0.1 for models with larger image sizes than 128×128 .

We observe that only ~8 denoising steps are sufficient for near-optimal performance, with modest degradation beyond that. Typically, diffusion models using the basic DDPM [30] sampler require many

hundreds of steps for decent results, but discrete data may have lowered it. It is unclear to us why the model performance degrades with more steps. Further research is needed and perhaps there is some performance to be gained by preventing this collapse.

We find similar classifier-free guidance values as those commonly used for text-to-image models. A guidance weight around 1 seems to help conditioning without collapsing diversity. We only explored image guidance, but future work could extend to other promptable signals such as weak supervision (points, boxes, scribbles), class labels, or few-shot examples. Modern universal segmentation systems must be promptable to be useful in practical settings. The ability to control condition strength on these inputs would provide a whole new dimension to promptable segmentation that traditional non-diffusion models do not have.

Overall, we provide a concrete path to make diffusion models viable for universal segmentation: analog bit diffusion for discrete labels, a noise schedule with input scaling, LAP for agnostic supervision, and a robust loss weighting. These choices yield consistent gains and make the method competitive in agnostic/holistic settings. At the same time, in broad foundation scenarios dominated by mask-classification architectures, our current model does not yet surpass strong discriminative baselines such as MaskFormer/Mask2Former or promptable SAM variants [2, 9, 10]. This gap likely reflects scale (data, compute, pretraining) and it motivates future work based on pretrained networks.

6 Conclusion

Diffusion models can serve as a viable framework for universal segmentation when adapted to discrete labels. It is necessary to modify the model to suit the discrete domain. Analog bits prove to be an effective encoding scheme, combined with a 2D gray code location-aware palette. Other effective modifications are an input-scaled noise schedule, x-prediction and using tanh as a final activation function.

While our approach does not yet surpass leading mask-based universal models in general foundation settings [2, 9, 10], it narrows the gap and offers capabilities those models lack: principled ambiguity modeling and sample-based exploration of plausible masks. Given the progress of diffusion segmenters like UniGS [7], we see a possible path forward: combine large-scale pretraining with analog bits and as many as the other proposed model improvements. Another path which might be more useful in practice is to integrate promptable conditioning combined with classifier-free guidance. If successful, generative universal segmenters could prove to be competitive models that remain both agnostic and holistic.

527

528

529

530

553

557

559

560

561

562

563

565

566

567

568

571

575

581

582

591

595

References

- K. He, G. Gkioxari, P. Dollár, and R. Girshick. 494 "Mask R-CNN". In: *PAMI* (2017). DOI: 10. 495 48550/ARXIV.1703.06870. URL: https:// 496 arxiv.org/abs/1703.06870. 497
- A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. 498 Rolland, L. Gustafson, T. Xiao, S. White-499 head, A. C. Berg, W.-Y. Lo, P. Dollár, 500 and R. Girshick. "Segment Anything". In: 501 arXiv:2304.02643 (2023). 502
- N. Ravi, V. Gabeur, Y.-T. Hu, R. Hu, C. 503 Ryali, T. Ma, H. Khedr, R. Rädle, C. Rol-504 land, L. Gustafson, E. Mintun, J. Pan, K. V. 505 Alwala, N. Carion, C.-Y. Wu, R. Girshick, P. 506 Dollár, and C. Feichtenhofer. "SAM 2: Segment Anything in Images and Videos". In: 508 arXiv preprint arXiv:2408.00714 (2024). URL: 509 https://arxiv.org/abs/2408.00714. 510
- R. Rombach, A. Blattmann, D. Lorenz, P. 511 Esser, and B. Ommer. "High-Resolution Im-512 age Synthesis with Latent Diffusion Models". 513 In: CoRR abs/2112.10752 (2021). arXiv: 2112. 514 10752. URL: https://arxiv.org/abs/2112. 515 10752. 516
- T. Chen. On the Importance of Noise Schedul-517 ing for Diffusion Models. 2023. arXiv: 2301. 518 10972 [cs.CV]. URL: https://arxiv.org/ 519 520 abs/2301.10972.
- T. Chen, R. Zhang, and G. Hinton. "Ana-521 log Bits: Generating Discrete Data using Diffusion Models with Self-Conditioning". In: 523 ICLR (2023). arXiv: 2208.04202 [cs.CV]. 524 URL: https://arxiv.org/abs/2208.04202. 525
 - L. Qi, L. Yang, W. Guo, Y. Xu, B. Du, V. Jampani, and M.-H. Yang. UniGS: Unified Representation for Image Generation and Segmentation. 2023. arXiv: 2312.01985 [cs.CV]. URL: https://arxiv.org/abs/2312.01985.
- N. Carion, F. Massa, G. Synnaeve, N. Usunier, 531 A. Kirillov, and S. Zagoruyko. "End-to-End 532 Object Detection with Transformers". In: 533 ECCV (2020). DOI: 10.48550/ARXIV.2005. 12872. URL: https://arxiv.org/abs/2005. 535 12872. 536
- B. Cheng, A. G. Schwing, and A. Kirillov. 537 "Per-Pixel Classification is Not All You Need for 538 Semantic Segmentation". In: NeurIPS (2021). 539 DOI: 10.48550/ARXIV.2107.06278. URL: 540 541 https://arxiv.org/abs/2107.06278.
- [10]B. Cheng, I. Misra, A. G. Schwing, A. Kir-542 illov, and R. Girdhar. "Masked-attention Mask 543 Transformer for Universal Image Segmenta-544 tion". In: CVPR (2022). DOI: 10.48550/ 545 ARXIV. 2112. 01527. URL: https://arxiv. 546 org/abs/2112.01527.

- J. Jain, J. Li, M. Chiu, A. Hassani, N. Orlov, 548 and H. Shi. "OneFormer: One Transformer to Rule Universal Image Segmentation". In: 550 CVPR (2023). DOI: 10.48550/ARXIV.2211. 551 06220. URL: https://arxiv.org/abs/2211. 552 06220.
- S. A. A. Kohl, B. Romera-Paredes, C. Meyer, 554 J. D. Fauw, J. R. Ledsam, K. H. Maier-Hein, S. M. A. Eslami, D. J. Rezende, and O. Ronneberger. "A Probabilistic U-Net for Segmentation of Ambiguous Images". In: Advances in Neural Information Processing Systems (NeurIPS) (2018).
- I. Bhat, J. P. W. Pluim, and H. J. Kuijf. "Generalized Probabilistic U-Net for medical image segementation". In: Uncertainty for Safe Utilization of Machine Learning in Medical Imaging (UNSURE) (2022). arXiv: 2207.12872 [cs.CV]. URL: https://arxiv.org/abs/ 2207.12872.
- I. Bhat, J. P. Pluim, M. A. Viergever, and H. J. Kuijf. "Effect of latent space distribution on the segmentation of images with multiple annotations". In: Machine Learning for Biomedical Imaging (MELBA) 2 (Apr. 2023), pp. 151–171. ISSN: 2766-905X. DOI: 10.59275/ j.melba.2023-18ae. URL: http://dx.doi. 574 org/10.59275/j.melba.2023-18ae.
- C. F. Baumgartner, K. C. Tezcan, K. Chai- 576 |15|tanya, A. M. Hötker, U. J. Muehlematter, K. Schawkat, A. S. Becker, O. Donati, and E. 578 Konukoglu. "PHiSeg: Capturing Uncertainty in Medical Image Segmentation". In: International Conference on Medical Image Computing and Computer-Assisted Intervention (MIC-CAI) (2019). arXiv: 1906.04045 [eess.IV]. 583 URL: https://arxiv.org/abs/1906.04045.
- K. Zepf, S. Wanna, M. Miani, J. Moore, J. 585 [16]Frellsen, S. Hauberg, F. Warburg, and A. Fer- 586 agen. "Laplacian Segmentation Networks Im- 587 prove Epistemic Uncertainty Quantification". In: International Conference on Medical Image Computing and Computer-Assisted Intervention. Springer Nature Switzerland Cham. 2024, pp. 349-359.
- T. Amit, E. Nachmani, T. Shaharbany, and L. 593 Wolf. "Segdiff: Image segmentation with diffusion probabilistic models". In: arXiv preprint arXiv:2112.00390 (2021).
- J. Wolleb, R. Sandkühler, F. Bieder, P. Val- 597 [18]maggia, and P. C. Cattin. "Diffusion Models for Implicit Image Segmentation Ensembles". In: Proceedings of The 5th International Conference on Medical Imaging with Deep Learning. Ed. by E. Konukoglu, B. Menze, A. Venkataraman, C. Baumgartner, Q. Dou, 603

- and S. Albarqouni. Vol. 172. Proceedings of Machine Learning Research. PMLR, 2022, pp. 1336-1348. URL: https://proceedings. mlr.press/v172/wolleb22a.html.
- [19] J. Wu, H. Fang, Y. Zhang, Y. Yang, and Y.
 Xu. "MedSegDiff: Medical Image Segmentation with Diffusion Probabilistic Model". In:
 arXiv preprint arXiv:2211.00611 (2022).
- [20] J. Wu, R. Fu, H. Fang, Y. Zhang, and Y.
 Xu. "MedSegDiff-V2: Diffusion based Medical Image Segmentation with Transformer".
 In: arXiv preprint arXiv:2301.11798 (2023).
- [21] T. Chen, L. Li, S. Saxena, G. Hinton, and D. J.
 Fleet. "A generalist framework for panoptic segmentation of images and videos". In: arXiv preprint arXiv:2210.06366 (2022).
- [22] T. Salimans and J. Ho. "Progressive Distillation for Fast Sampling of Diffusion Models".
 In: The International Conference on Learning Representations (ICLR) (2022).
- [23] D. P. Kingma, T. Salimans, B. Poole, and J.
 Ho. Variational Diffusion Models. 2023. arXiv:
 2107.00630 [cs.LG]. URL: https://arxiv.org/abs/2107.00630.
- E. Hoogeboom, J. Heek, and T. Salimans. Simple diffusion: End-to-end diffusion for high resolution images. 2023. arXiv: 2301.11093 [cs.CV]. URL: https://arxiv.org/abs/2301.11093.
- E. Hoogeboom, T. Mensink, J. Heek, K. Lamerigts, R. Gao, and T. Salimans. Simpler Diffusion (SiD2): 1.5 FID on ImageNet512 with pixel-space diffusion. 2024. arXiv: 2410. 19324 [cs.CV]. URL: https://arxiv.org/abs/2410.19324.
- [26] T. Strang, A. Dammann, M. Röckl, and S. 639 Plass. "Using Gray codes as Location Iden-640 tifiers". In: 6. GI/ITG KuVS Fachgespräch 641 Ortsbezogene Anwendungen und Dienste. PDF 642 available at DLR's elib repository. Institut 643 für Kommunikation und Navigation, German 644 Aerospace Center (DLR). Oberpfaffenhofen, Germany, Oct. 2009. 646
- [27] Q. Lu, J. Kuen, S. Tiancheng, G. Jiuxiang, G.
 Weidong, J. Jiaya, L. Zhe, and Y. Ming-Hsuan.
 "High-Quality Entity Segmentation". In: ICCV.
 2023.
- [28] I. Loshchilov and F. Hutter. "Decoupled Weight Decay Regularization". In: International Conference on Learning Representations. 2019. URL: https://openreview.net/forum?id=rk6qdGgCZ.
- [29] J. Ho and T. Salimans. "Classifier-Free Diffusion Guidance". In: NeurIPS (2021). arXiv:
 2207.12598 [cs.LG].

DJ. Ho, A. Jain, and P. Abbeel. "Denoising 659 Diffusion Probabilistic Models". In: *CoRR* 660 abs/2006.11239 (2020). arXiv: 2006.11239. 661 URL: https://arxiv.org/abs/2006.11239. 662