

# Uncertainty-Aware ControlNet: Bridging Domain Gaps with Synthetic Image Generation

Joshua Niemeijer<sup>1</sup> Jan Ehrhardt<sup>2,3</sup> Heinz Handels<sup>2,3</sup> Hristina Uzunova<sup>3</sup>

<sup>1</sup>German Aerospace Center (DLR) <sup>2</sup>University of Lübeck

<sup>3</sup>German Research Center for Artificial Intelligence (DFKI)

Joshua.Niemeijer@dlr.de, {jan.ehrhardt, heinz.handels}@uni-luebeck.de,  
hristina.uzunova@dfki.de

## Abstract

*Generative Models are a valuable tool for the controlled creation of high-quality image data. Controlled diffusion models like the ControlNet have allowed the creation of labeled distributions. Such synthetic datasets can augment the original training distribution when discriminative models, like semantic segmentation, are trained. However, this augmentation effect is limited since ControlNets tend to reproduce the original training distribution.*

*This work introduces a method to utilize data from unlabeled domains to train ControlNets by introducing the concept of uncertainty into the control mechanism. The uncertainty indicates that a given image was not part of the training distribution of a downstream task, e.g., segmentation. Thus, two types of control are engaged in the final network: an uncertainty control from an unlabeled dataset and a semantic control from the labeled dataset. The resulting ControlNet allows us to create annotated data with high uncertainty from the target domain, i.e., synthetic data from the unlabeled distribution with labels. In our scenario, we consider retinal OCTs, where typically high-quality Spectralis images are available with given ground truth segmentations, enabling the training of segmentation networks. The recent development in Home-OCT devices, however, yields retinal OCTs with lower quality and a large domain shift, such that out-of-the-pocket segmentation networks cannot be applied for this type of data. Synthesizing annotated images from the Home-OCT domain using the proposed approach closes this gap and leads to significantly improved segmentation results without adding any further supervision. The advantage of uncertainty-guidance becomes obvious when compared to style transfer: it enables arbitrary domain shifts without any strict learning of an image style. This is also demonstrated in a traffic scene experiment. The implementation is available: <https://github.com/JNiemeijer/UnACorN.git>*

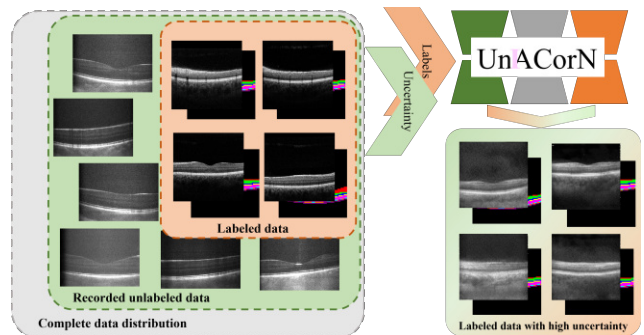


Figure 1. We aim at learning the uncertainties of a recorded but unlabeled distribution and generating new labeled data containing high uncertainties.

## 1. Introduction

With the shift of research interest towards deep neural network-based approaches, the problem of labeled data scarcity becomes increasingly prominent. Typically, discriminative models, like semantic segmentation networks, require large and representative training distributions with given ground truth annotations. However, such distributions are inherently hard to create. Firstly, manual labeling effort is costly and time-consuming. Secondly, the test data might expand to previously unknown points of the distribution, making the creation of a realistic representative training dataset impossible. This problem becomes even further aggravated when it comes to the medical image domain since data availability is generally limited, rare diseases lead to unexpected distributions, and privacy concerns need to be taken into account. Furthermore, most labeling tasks require expert knowledge about the specific domain. Thus, existing pre-trained methods are often ill-suited to newly acquired datasets (e.g., using a different device or a novel acquisition technique). A known technique to work against data scarcity, in general, is data augmentation. The most common data augmentation type includes simple im-

age transformations like affine/elastic image deformations and adding noise to the image intensities. Such augmentations typically extend the training data distribution and, if chosen carefully, lead to better performance. However, with this simple type of augmentation, most cases would conform to the training data domain and would not perform well with large domain shifts in the test dataset. A more sophisticated way of data augmentation is enabled by generative methods, where new data domains can be synthetically produced. In order to enable the usage of the generated data for training, the data is commonly generated conditioned on given ground truth labels. Prominent examples of unpaired domain translation methods are CycleGANs [24], which can translate two unlabeled and unpaired distributions into each other and generate realistic images of a new domain. However, CycleGAN-generated images often do not conform to the shape and, thus, semantic labels of the source domain and cannot directly be applied for training. Ideas to overcome this problem are presented by approaches like CyCADA [3], or ACCUT [16]. However, these methods require a reliable segmentation approach for the target data domain, which might be challenging depending on the application. Furthermore, there is a variety of unsupervised topology-preserving domain translation approaches specifically designed for medical images (e.g., [19, 20]).

Naturally, led by the current state-of-the-art in image generation, diffusion models that stand out for their exceptional capacity to produce realistic images and offer excellent control possibilities like ControlNet [6, 18, 22, 23], approaches like MAISI [2] have emerged for large-scale image generation in the medical image domain. Even though ControlNets allow for the creation of high-quality labeled distributions by introducing image labels as a control mechanism, generative models like GANs and diffusion models are known to reproduce their training data distribution. This essentially limits their augmentation effect since unknown distribution parts that might be present in the test dataset will most likely not be synthesized. Instead of selective control of the generated image distribution, we want to shift the paradigm towards an automated control of the synthesis process in order to generate images that do not conform to the predetermined distribution of the generative model. However, we are regulated by explicit requirements of the downstream task.

We pursue the idea of generating labeled unseen image distributions, specifically targeting the relevance of the data for a downstream task. This requires a metric that enables a statement about data relevance. Inspired by methods of active learning [7], where relevant images are selected for training based on epistemic uncertainty, we propose to utilize this metric in a generative manner. Previous works like TSynD [9], where the generation process is steered towards data with high epistemic uncertainty, have shown the

efficacy of this consideration. However, TSynD requires a time- and memory-consuming optimization of the latent space of a generative model and is only able to slightly adapt an input image to the given uncertainty level, offering neither a possibility to generate completely new data nor a guarantee of topology preservation of the input data.

In this work, we introduce **UnACorN**: an **uncertainty-aware ControlNet**, which enables the generation of labeled image data with high epistemic uncertainty towards a given discriminative model. The proposed method combines two independent conditioning strategies: 1) semantic conditioning on the image labels, trained on a given labeled data distribution, and 2) conditioning on the uncertainty, trained on an unlabeled image domain. In this way, unlabeled images can be utilized for training, as well, where their uncertainty for a particular segmentation model is mapped to image information, and a labeled data distribution with control over uncertainty levels can be produced. The main contributions of our work are:

- We propose an *uncertainty-based conditioning mechanism* for ControlNets, enabling the training on existing unlabeled datasets with large domain gaps without the need for any semantic control information.
- Unlike the conventional ControlNet, where only data from the labeled distribution can be generated, we enable the combination of a semantic label control and uncertainty control, leading to the synthesis of *diverse labeled datasets*.
- In contrast to traditional style-transfer methods, bound to the statistical distribution of the underlying data domains, our approach offers the possibility to *sample data based on different uncertainty levels* and, thus, generate datasets covering larger parts of the distribution.

In our experiments, we evaluate UnACorN on a distribution shift scenario in retinal OCTs, where we have a labeled distribution of the commonly acquired Spectralis OCTs and an unlabeled distribution of the novel self-examination technique HOME-OCTs [21], that are structurally different from each other (examples in Fig. 7) [8, 15, 16]. Hence, we learn control over the class on the Spectralis dataset and control over uncertainty on both. In our experiments, we demonstrate that our approach to data generation allows for improved training when compared to common style transfer approaches. However, UnACorN is not yet another domain translation method: its ability to generate data based on uncertainty enables the representation of arbitrary domain shifts without learning a particular image style, which we illustratively show for traffic scene images.

## 2. Methods

In this work, we present an "Uncertainty-Aware ControlNet": UnACorN. Our goal is to create training data distributions for a discriminative model, in this case, a segmentation

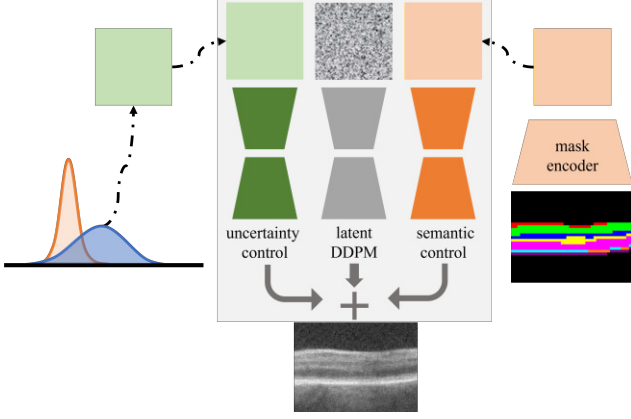


Figure 2. During inference, we infer both the Semantic-ControlNet and the Uncertainty-ControlNet in each diffusion step. Finally, a weighted sum of both is computed, allowing for control over the segmentation mask and the uncertainty in the image.

neural network. The distributions should have the following properties: the data should match given ground-truth labels (“labeled”), and they should be highly tailored to the particular task (“relevant”).

In order to fulfill the first property, i.e., “labeled”, we utilize the ControlNet [23] approach on a labeled dataset coming from a distribution  $(x, y) \sim p(\mathcal{X}^{\mathcal{L}}, \mathcal{Y})$ , with  $\mathcal{X}^{\mathcal{L}}$  denoting images acquired with specific devices or parameters and  $\mathcal{Y}$  are label maps of the underlying anatomy or scene. The resulting Semantic-ControlNet can create images  $\hat{x}$  that correspond to a given label map. However, the generated images will correspond to the labeled data distribution  $\hat{x} \sim p(\mathcal{X}^{\mathcal{L}})$  (see Fig. 4). Training on such data, therefore, does not provide additional information to the labeled data.

To fulfill the second property, i.e., “relevant”, we need to condition the generated images to a certain appearance. A common approach is to use active learning to select data that is particularly relevant to a task and requires labeling. So-called acquisition functions are utilized to select images with a high epistemic uncertainty for manual labeling. The new labeled data can then serve as additional input to a ControlNet. In our approach, we utilize this idea in an alternative way, which does not require human intervention for labeling. We propose to employ the (epistemic-) uncertainty of a downstream discriminative model as an indicator function for identifying relevant parts of the unlabeled distribution. Therefore, in our application, we train a second separate ControlNet conditioned on the epistemic uncertainty of a segmentation model trained on the labeled distribution. Since we can determine the uncertainty in an unsupervised way, we can train this Uncertainty-ControlNet on unlabeled but recorded (known) distributions  $x \sim p(\mathcal{X}^{\mathcal{U}})$ . The resulting Uncertainty-ControlNet should, therefore, be able

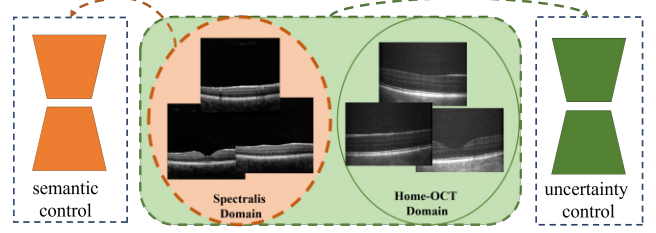


Figure 3. The Semantic-ControlNet is trained on the labeled Spectralis Domain. The Uncertainty-ControlNet is trained on both the Spectralis and Home-OCT domains.

to create “known-unknown” information, i.e., image information that was recorded but is not labeled and is difficult for the segmentation network to process.

Given the two models, the Semantic-ControlNet and the Uncertainty-ControlNet, we have control over the two desired properties: 1. labels, and 2. relevance. In order to create the image, we fuse the predictions in each diffusion step as shown in Fig. 2.

## 2.1. Pre-Training

Our approach consists of three building blocks: a denoising diffusion probabilistic model (DDPM), the Semantic-ControlNet, and the Uncertainty-ControlNet.

**DDPM pre-training:** We assume/ train a DDPM with high expressiveness to create a wide variety of images, as is common for current foundation models. In particular, the DDPM must be able to generate images from both distributions  $\mathcal{X}^{\mathcal{L}}$  and  $\mathcal{X}^{\mathcal{U}}$ , i.e., notably images that were not used during the training of the downstream task. Pre-training or fine-tuning of the DDPM is therefore only necessary if the use case involves new acquisition devices or rare imaging conditions, and can then be performed using unlabeled data. Such a case is given for our medical use-case; for the automotive use-case, we can re-utilize existing foundation models. In our approach, we employ latent diffusion models (LDMs) [12], i.e., all images are transformed into a low-dimensional latent space representation  $z \in \mathcal{Z}$  by a pre-trained encoder-decoder model. The DDPM predicts the noise in each denoising time-step  $t$ :

$$\epsilon(z_t, t) = \mathcal{LDM}(z_t, t), \quad (1)$$

where  $z_t$  is a noisy version of the encoded latent space representation  $z = \mathcal{E}(x)$  of image  $x$ , where  $\mathcal{E}$  denotes the encoder.

**Semantic-ControlNet:** As shown in Fig. 3, the Semantic-ControlNet is trained on the labeled data  $(x, y) \sim p(\mathcal{X}^{\mathcal{L}}, \mathcal{Y})$  only. As Fig. 4 shows, it, therefore, reproduces the labeled distribution during inference. We denote the semantic control network as  $\mathcal{C}_S$ , which takes as input the current latent space  $z_t$  and the time step  $t$  as well as an encoded

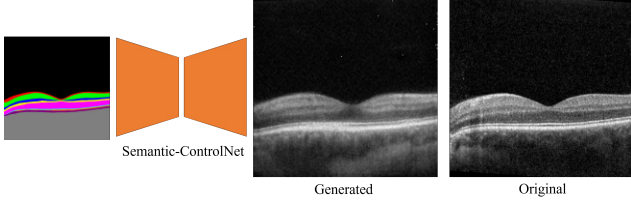


Figure 4. The images generated by the Semantic-ControlNet match the distribution of the labeled images.

control image  $c_s = \mathcal{E}_S(y)$  representing the segmentation mask. The network  $\mathcal{C}_S$  generates a noise prediction:

$$\epsilon_S(z_t, c_s, t) = \mathcal{C}_S(z_t, c_s, t). \quad (2)$$

The encoder  $\mathcal{E}_S$  is trained along with the ControlNet. Its task is to adapt the spatial size of the label map to the size of the latent space without introducing artifacts, as would occur with plain down-sampling.

**Uncertainty-ControlNet:** A segmentation network  $\hat{y} = \mathcal{S}(x)$  was pre-trained based on labeled data  $(x, y) \sim p(\mathcal{X}^L, \mathcal{Y})$ . We utilize the uncertainty of the pre-trained segmentation network  $\mathcal{S}$  as a condition for the training of the Uncertainty-ControlNet. As shown in Fig. 3, uncertainty measurements of both image distributions, labeled and unlabeled images, are used for training. Our measure of uncertainty is the pixelwise entropy:

$$H(p(y|x_{i,j})) = - \sum_{y \in \mathcal{Y}} \hat{p}(y|x_{i,j}) \log(\hat{p}(y|x_{i,j})). \quad (3)$$

The pixelwise entropy itself contains a lot of information about the spatial structure of the image. In our use case, this is not a desirable feature, as the Semantic-ControlNet generates the image structure, while the uncertainty describes the confidence of the segmentation network with the image content. We, therefore, compute the normalized mean uncertainty of the image and scale it to the range  $[0, 100]$ :

$$\mathbf{U}_H(x) = \frac{100}{N \cdot \log(L)} \sum_{i,j} H(p(y|x_{i,j})), \quad (4)$$

where  $N$  is the number of image pixels, and  $L$  is the number of labels in the segmentation mask.

The resulting uncertainty value is expanded to a control image where each pixel has the value of the average uncertainty. This control image serves as input to the Uncertainty-ControlNet. During inference, we can then sample from the uncertainty distributions. Fig. 5 shows the different distributions of average image uncertainty of the labeled and unlabeled datasets. As visualized, the images generated corresponding to given uncertainty values shift from the style of the labeled distribution to the style of the

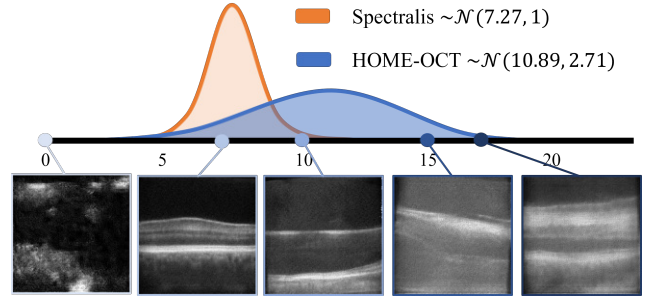


Figure 5. The normal distributions of the average image uncertainties of the segmentation network applied to the Spectralis and Home OCT data. We sample from these distributions and feed the resulting uncertainty to the Uncertainty-ControlNet. The response to different uncertainties can be seen in the bottom row of images.

unlabeled distribution with increasing uncertainty. We denote the Uncertainty-ControlNet as  $\mathcal{C}_U$ . The  $\mathcal{C}_U$  network takes a latent representation  $z_t$ , the control image  $c_u$  representing the average uncertainty, and a time step  $t$  as inputs, and produces a noise prediction:

$$\epsilon_U(z_t, c_u, t) = \mathcal{C}_U(z_t, c_u, t). \quad (5)$$

## 2.2. Inference

Fig. 2 shows the final architecture that is used during inference. We utilize the segmentation labels from the labeled distribution as input for the Semantic-ControlNet. At the same time, we infer the Uncertainty-ControlNet.

**Sampling:** The uncertainties are sampled from a Gaussian distribution corresponding to the uncertainty distribution of the unlabeled data. To create this Gaussian, we first apply our segmentation network to the images from the unlabeled distribution and determine the mean and standard deviation. During sampling, we then randomly sample from the resulting Gaussian and create the uncertainty-control image from it. Therefore, the Uncertainty-ControlNet is steering the generation process to create images from the unseen distribution in each diffusion step.

**Weighting:** During the parallel inference of the Semantic-ControlNet and the Uncertainty-ControlNet, we fuse the down and mid-block noise prediction by computing a weighted sum. Let  $\epsilon_U(z_t, c_u, t)$  and  $\epsilon_S(z_t, c_s, t)$  denote the noise predictions from the Uncertainty-ControlNet and the Semantic-ControlNet, respectively, for an input  $x$  at time step  $t$ . The fused noise prediction is given by

$$\epsilon_{\text{fused}}(z_t, c_u, c_s, t) = \alpha \epsilon_U(z_t, c_u, t) + \epsilon_S(z_t, c_s, t) \quad (6)$$

where  $\alpha \in [0, 1]$  controls the contribution of the Uncertainty-ControlNet.  $\alpha$  is empirically determined. The LDM and the fused noise prediction of the ControlNets define the final noise prediction:

$$\epsilon(z_t, t, c_s, c_u) = \epsilon_{\text{fused}}(z_t, c_s, c_u, t) \oplus \mathcal{LDM}(z_t, t), \quad (7)$$



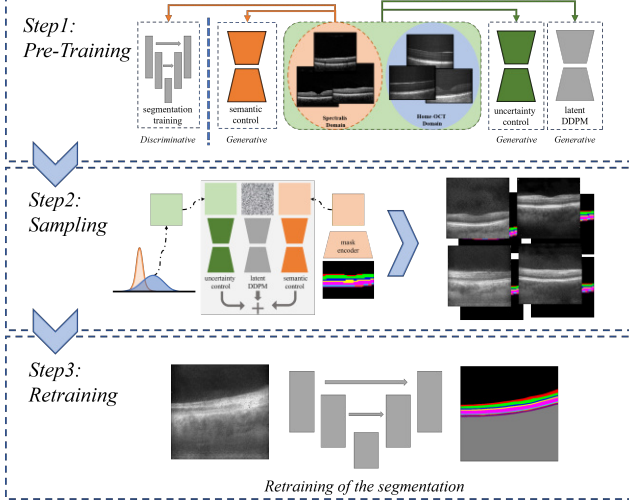


Figure 6. Schematic depiction of the three steps of our approach. Step 1: The generative blocks and the segmentation model are trained. Step 2: The new training distribution is sampled. Step 3: The segmentation network is retrained.

where  $\oplus$  is a  $1 \times 1$  convolution.

### 2.3. Assembling the Bricks

The final architecture of the UnACorN approach consists of three main steps, as shown in Fig. 6. First, we pre-train the diffusion blocks and the segmentation model. The diffusion blocks are the backbone LDM, the Semantic-ControlNet, and the Uncertainty-ControlNet. During the training of the latter, the frozen pre-trained segmentation model on the labeled images provides the uncertainty control for all images.

The second step is the data generation using the pre-trained generative blocks as shown in Fig. 2. For this, uncertainty values from a normal distribution parametrized with the mean and standard deviation based on the real HOME-OCT images are sampled and input to the Uncertainty-ControlNet, while a label map from a Spectralis image is used as input to the Semantic-ControlNet. This way, a new labeled dataset with high uncertainties is generated.

Finally, in the third step, we retrain our discriminative network on the generated data. This process reduces the uncertainties from the model and yields improved segmentation results for out-of-distribution data.

## 3. Experiments and Results

In our experiments, we mainly pursue a medical image processing scenario based on retinal OCTs. Here, we consider two structurally different distributions of OCT-data coming from different sensor devices: Spectralis and HOME-OCT, generally assuming the Spectralis to represent the labeled data distribution and HOME-OCT to be unlabeled.

We show that we can create a labeled distribution based on the segmentations from the Spectralis domain, containing the uncertainties observed in the HOME-OCT domain, and compare the results to style transfer approaches. Additionally, we show an application to traffic scenes to prove the generalization ability of our method.

### 3.1. Data

**Retinal OCTs:** In our experiments, an in-house OCT dataset from 50 healthy subjects is used [17]. For each subject, both eyes are recorded with the Spectralis OCT device (Heidelberg Engineering, Heidelberg, Germany) and a prototype of a self-examination low-cost full-field (HOME) OCT device, yielding two different OCT image domains (Fig. 7). As typical for retinal OCTs, 2D B-Scans are extracted from the images and resampled to  $256 \times 256$  pixels, resulting in  $\sim 10000$  images per domain. Both domains contain ground-truth segmentations of the anatomical retinal layers; however, the HOME-OCT labels are only considered for evaluation in our experiments.

**Street Scenes:** The Cityscapes [1] dataset (DCS) contains images taken from German cities under a stable sensor setup in very consistent environmental conditions. The ACDC [14] dataset is taken in Zürich in different weather conditions. It contains rainy, snowy, foggy, and night scenes. Since the DCS only contains day images with few lighting/weather conditions, there is a large domain gap between the datasets. Both of these datasets provide semantic segmentation annotations. We utilize DCS as our labeled distribution and ACDC as our unlabeled distribution.

### 3.2. Uncertainty-guided Retinal OCT Generation

The primary objective of this experiment is to synthesize data that induces high uncertainty in a given pre-trained retinal layer segmentation model. The synthetic data is leveraged to refine the model through retraining, enhancing segmentation performance on previously unseen image domains. In our experiments, we consider two image domains: a labeled source domain (Spectralis) and an unlabeled target domain (HOME-OCT).

#### 3.2.1. Models

Our method consists of three main building blocks: The backbone DDPM, the Semantic-ControlNet, and the Uncertainty-ControlNet.

**Backbone DDPM:** For the synthesis of data, we assume that a generative model pre-trained on multiple domains exists. Due to the specificity of the retinal OCT data, here, a latent diffusion model based on [11] was trained from scratch with the given datasets. However, with the current shift of research focus towards foundation models, this step may easily be omitted in the future. We plan on releasing the trained diffusion model.

**Semantic-ControlNet:** We train the Semantic-ControlNet on the Spectralis domain for 150 epochs with a learning rate of  $2.5e-5$ .

**Uncertainty-ControlNet:** The Uncertainty-ControlNet is trained on the combined unlabeled and labeled distributions. As described in section 2, we use the inference of a pre-trained segmentation network to compute the average entropy over the pixels according to Eq. 4. The average entropy is assigned to every pixel of the control image that serves as input to the Uncertainty-ControlNet. This training is done for 150 epochs with a learning rate of  $2.5e-5$ .

Furthermore, for the uncertainty calculation, a retinal segmentation network is required. Here, we train a U-Net-architecture [13] with Dice-loss for a maximum of 400 epochs using early stopping and a learning rate of  $1e-4$ . For uncertainty estimations, we pre-train the retinal layer segmentation neural network on the Spectralis data only and use the network’s inference to calculate the uncertainty on both image domains.

### 3.2.2. Data Generation

During inference, we infer the diffusion process with 1000 time-steps. In each time step, the noise prediction of the Semantic-ControlNet and the Uncertainty-ControlNet are fused by computing a weighted sum. We empirically found the weight factor  $\alpha = 0.4$  for the noise prediction of the Uncertainty-ControlNet to work well. For each of the training images of the labeled distribution, we create a synthetic image. Generally speaking, our method would allow for the creation of larger distributions. However, to ensure fairness in comparison with the style transfer approaches used here, we create the same number of images. To introduce more diversity into the distribution, we sample from the distribution of average image uncertainties from the unlabeled domain. This distribution is modeled as a Gaussian  $\mathcal{N}(10.89, 2.71)$  as shown in Fig. 5, where the mean and standard deviation are determined for the unlabeled data. Note that higher uncertainty values correspond to images that would be more challenging for segmentation (noisier, less visible layers), whereas lower uncertainty values, closer to the mean of the Spectralis distribution, yield images similar to the real Spectralis data.

### 3.2.3. Results

As a surrogate of data quality in our experimental setup, we use generated and real datasets to train a retinal layer segmentation U-Net and observe its segmentation accuracy behavior. For comparison, we chose two common generative approaches: cycleGAN [24] and CUT [10], which are trained to generate domain-translated data (Spectralis→HOME-OCT) in an unsupervised fashion. We generate training datasets of a rigid size ( $\sim 1100$  images) using the different approaches and evaluate on a rigid size test set ( $\sim 1100$  images) consisting of real HOME-OCTs.

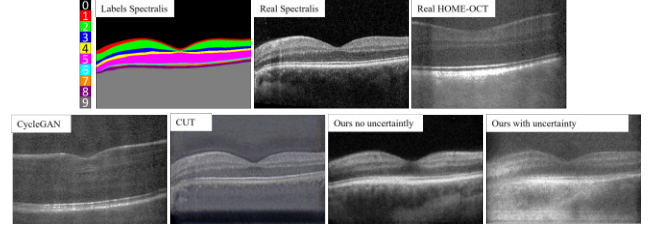


Figure 7. Example image data of the real Spectralis and HOME-OCT images (top row); and synthesized images (bottom row), from left to right: generated using CycleGAN; generated using CUT; generated using only the Semantic-ControlNet; generated using the semantically- and uncertainty-conditioned ControlNets.

Overall, the following training setups are evaluated: real HOME-OCT images (upper bound), real Spectralis OCTs (lower bound); synthetic images generated by CycleGAN, CUT, and the proposed UnACorN; synthetic image datasets combined with real Spectralis data. The results of this experiment can be seen in Tab. 1, where we show the mean intersection over union values (mIoU, also known as Jaccard index) over all retinal labels using the different training settings. Overall, the best training scenario features the dataset consisting of our images and the real Spectralis data, demonstrating the efficiency of our approach for data augmentation. Using the uncertainty-guided dataset generation, we achieve a mIoU of 0.65, which is a vast improvement compared to training on the real Spectralis data with a large domain gap.

The results also show that the CycleGAN-generated training dataset is the worst-performing in this use case. This is because CycleGAN does not offer any semantic consistency mechanism, leading to realistic synthetic HOME-OCT images that do not align with the shape and labels of the corresponding Spectralis images. Thus, an inaccurately labeled data distribution is generated (example in Fig. 7). CUT, on the other hand, offers unsupervised semantic consistency and results in label-conform data distributions, as evidenced by better segmentation results. However, as visible in Fig. 7, the established domain shift is not as strong as the one achieved by our method.

To evaluate the similarity between the domains of the generated data and the real HOME-OCT data, we compute the commonly used FID (Fréchet inception distance) metric, which measures the similarity between distributions (Tab. 2). These results emphasize that CycleGAN indeed establishes a successful domain adaptation, as it is, per design, a style-transfer approach. However, due to the large spatial bias in the Spectralis and HOME-OCT domains, CycleGAN is simply able to map the distribution of the HOME-OCT data containing very large spatial shifts w.r.t. the given retinal labels. Furthermore, the large FID values of the semantically constrained methods, CUT and Un

lbl	real data		synthetic data					
	HOME	Spect	Cycle CUT ours			Cycle CUT ours		
	upper bound	lower bound				+real Spectralis		
0	0.99	0.86	0.93	0.85	0.83	0.92	0.90	<b>0.96</b>
1	0.71	0.33	0.45	0.41	0.35	0.49	<b>0.53</b>	0.51
2	0.86	0.51	0.65	0.58	0.60	0.75	0.73	<b>0.78</b>
3	0.69	0.37	0.47	0.52	0.41	0.59	0.33	<b>0.64</b>
4	0.61	0.41	0.35	0.40	0.21	0.50	0.50	<b>0.60</b>
5	0.86	0.65	0.41	0.63	0.65	0.60	0.74	<b>0.82</b>
6	0.67	0.28	0.13	0.17	0.37	0.19	0.36	<b>0.40</b>
7	0.60	0.36	0.01	0.25	<b>0.44</b>	0.31	0.38	0.35
8	0.74	0.61	0.10	0.51	<b>0.69</b>	0.58	0.62	0.60
9	0.96	0.84	0.48	0.80	0.77	0.74	0.81	<b>0.87</b>
$\emptyset$ all	0.77	0.52	0.40	0.51	0.53	0.57	0.59	<b>0.65</b>
$\emptyset[1 - 8]$	0.72	0.44	0.32	0.43	0.46	0.50	0.52	<b>0.59</b>

Table 1. Mean Jaccard index values of the retinal segmentation network, trained on different datasets: real HOME-OCT images (upper bound), real Spectralis OCTs (lower bound); synthetic images generated by CycleGAN, CUT, and the proposed UnACorN; synthetic image datasets combined with real Spectralis data. The last two rows show mean values over all labels and labels  $[1 - 8]$ , excluding the background labels.

real	real Spectralis	synthetic CycleGAN	synthetic CUT	synthetic UnACorN
HOME-OCT	152.6	68.9	173.5	132.5
Spectralis	-	151.3	155.4	151.9

Table 2. FID values for the created datasets compared to the real Spectralis and HOME-OCT data. Small values indicate similar distributions.

ACorN, show that the FID score is likely influenced by the spatial bias as opposed to only the image style of the distributions. Apart from that, the FID score shows that UnACorN is able to slightly reduce the gap between the domains by achieving  $FID = 132.5$  for the generated data, compared to  $FID = 152.6$  for the real Spectralis data. Also, the FID value between the real Spectralis and real HOME-OCT data distributions is very similar to the FID values of the generated data and the real Spectralis (152.6 vs. 151.9), again suggesting a good domain adaptation ability. Additionally, we measured the mean uncertainty values over the generated and real distributions as suggested in Eq. 4. The proposed UnACorN approach, as expected, yields a high mean epistemic uncertainty value of 12.85, which is significantly higher than the real Spectralis data (7.27) and similar to the mean over the real HOME-OCT data (10.89). Overall, we can deduce that the proposed UnACorN has created a labeled data distribution with a similar uncertainty level to the unlabeled HOME-OCT data without directly learning the style or domain properties of the latter.

### 3.3. Application to Street Scenes

Since no explicit knowledge about the domains is integrated into UnACorN, it becomes apparent that the proposed method is not yet another style transfer approach. UnACorN’s ability to generate data based on uncertainty enables the representation of arbitrary domain shifts without learning a particular image style, matching real-world scenarios where multiple (partly unknown) data domains are inhabiting a dataset. To show this generalization ability of our approach, we present a toy example on a street-scene use case, where our labeled distribution corresponds to the DCS data and the unlabeled distribution originates from the ACDC dataset. For this experiment, the backbone DDPM is a foundation model for RGB image generation [12], and the semantic ControlNet is based on the ALDM approach [5]. For the calculation of uncertainty, we utilize a SegFormer-based [4] architecture trained on the Cityscapes data. For data generation, we sample the uncertainty from  $\mathcal{N}(9.6, 2.4)$  and generate an uncertainty-based dataset of size 2975, which corresponds to generating a synthetic version of each Cityscapes train image (examples can be seen in 8). An architectural difference of this approach is, however, the fact that the stable diffusion model additionally allows for influencing the generation process with text prompts (Fig. 8). During the dataset creation, we include the four subdomains of the ACDC dataset (rain, fog, snow, and night) in the prompts, enabling targeted image generation for the baseline Semantic-ControlNet. The advantage of introducing uncertainty with our method, however, quickly becomes obvious. Since the control over the uncertainty in the image was trained on both the Cityscapes and the ACDC domain, the sub-domains of the unlabeled ACDC dataset, i.e., rain, snow, fog, or darkness, are generated more convincingly. This is because the Uncertainty-ControlNet learned the correlation of the uncertainties and the image information on the unlabeled data distribution. As Fig. 9 shows, this causes the generated images of our UnACorN approach to be closer in style to the original ACDC data than the Semantic-ControlNet trained on Cityscapes. To evaluate this effect quantitatively, we used the sampled dataset to train the SegFormer-based model by Hoyer et al. [4] on the generated synthetic data simultaneously with the original Cityscapes data. The resulting model was then evaluated on the ACDC domain. We achieved improvements of around 2% in mIoU compared to the model only trained on the Cityscapes dataset and an improvement of around 1% compared to the training on a dataset created with the same prompting strategy using the ALDM model.

## 4. Discussion and Conclusion

In this work, we proposed UnACorN: a generative approach for synthesizing labeled image data with high epis-





Figure 8. The left column shows the Semantic-ControlNet result by a model (ALDM) trained on the Cityscapes dataset. The right column shows the effect of the uncertainty guidance. Given the prompt, the image that is generated with UnACorN is visibly more “difficult”, making the segmentation network more uncertain (the fog scene gets foggier; the night scene becomes darker, etc.).

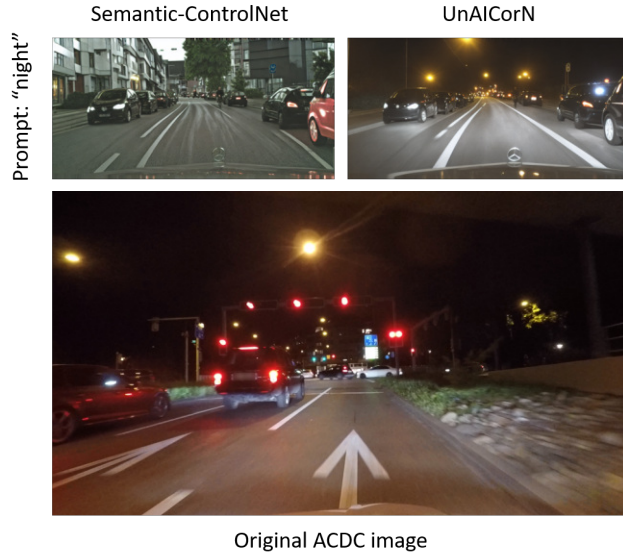


Figure 9. Comparison of an image generated by the baseline Semantic-ControlNet and our UnACorN approach with an original ACDC image of the subdomain from the prompt. Since the UnACorN has learned the correlation between uncertainty and image information on the ACDC dataset, the style matches better.

temic uncertainty towards a downstream task. We propose a novel uncertainty-based conditioning mechanism for ControlNets, which enables the modeling of the uncertainties of a pre-trained discriminative model observed on recorded but unlabeled data. Such uncertainties could be interpreted as “known unknowns”, since the image information that causes uncertain predictions is available; however, due to the lack of semantic labels, it cannot be utilized during the training process. In addition to uncertainty control, our method retains control over the semantic information, such as labels, during the generation process. In this way, both labeled and unlabeled data can be utilized during training, leading to an overall larger and more diverse training dataset and, eventually, to a more expressive generative model.

This dual control strategy enables the sampling of diverse labeled training datasets based on uncertainties. As our experiments show, datasets generated by this approach lead to improved segmentation results as compared to standard style-transfer methods. Such methods are rigid in their learning of the explicit domain information, whereas the proposed method does not require any explicit knowledge about the domain (e.g., the type of acquisition device used for medical images) since it offers uncertainty-guidance based on out-of-distribution examples. Future work may explore this property more in depth by purposely presenting diverse domain shifts in the unlabeled dataset.

An underexplored advantage of the presented approach is its ability to sample multiple instances per input label map, unlike most style transfer models. For each label map, uncertainties can be sampled from a Gaussian distribution, corresponding to the distribution of the unlabeled data domain, and create various versions of the same label. This property of the UnACorN approach was not applied during the comparison with the baseline methods to avoid an unfair advantage. In future work, however, we will explore this augmented generation, which is especially interesting for automotive use cases since car manufacturers usually have large amounts of unlabeled data from their fleets. Another interesting future work, UnACorN allows for the iterative improvement of the uncertainty modeling of the generative and discriminative models. After training, the uncertainties of the discriminative model change, and hence, the generative model can be updated. This circular dependency opens doors for further advancements.

## Acknowledgement

The research leading to these results is funded by the German Federal Ministry for Economic Affairs and Energy within the project “NXT GEN AI METHODS – Generative Methoden für Perzeption, Prädiktion und Planung”. The authors would like to thank the consortium for the successful cooperation.



## References

- [1] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The Cityscapes Dataset for Semantic Urban Scene Understanding. In *Proc. of CVPR*, pages 3213–3223, Las Vegas, NV, USA, 2016. 5
- [2] Pengfei Guo, Can Zhao, Dong Yang, Ziyue Xu, Vishwesh Nath, Yucheng Tang, Benjamin Simon, Mason Belue, Stephanie Harmon, Baris Turkbey, et al. Maisi: Medical ai for synthetic imaging. *arXiv preprint arXiv:2409.11169*, 2024. 2
- [3] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *Proceedings of the 35th International Conference on Machine Learning*, pages 1989–1998. PMLR, 2018. 2
- [4] Lukas Hoyer, Dengxin Dai, and Luc Van Gool. Daformer: Improving network architectures and training strategies for domain-adaptive semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 9914–9925. IEEE, 2022. 7
- [5] Yumeng Li, Margret Keuper, Dan Zhang, and Anna Khoreva. Adversarial supervision makes layout-to-image diffusion models thrive. In *The Twelfth International Conference on Learning Representations*, 2024. 7
- [6] Han Lin, Jaemin Cho, Abhay Zala, and Mohit Bansal. Ctrl-adapter: An efficient and versatile framework for adapting diverse controls to any diffusion model, 2024. 2
- [7] Sudhanshu Mittal, Joshua Niemeijer, Özgün Çiçek, Maxim Tatarchenko, Jan Ehrhardt, Jörg P Schäfer, Heinz Handels, and Thomas Brox. Realistic evaluation of deep active learning for image classification and semantic segmentation. *International Journal of Computer Vision*, pages 1–23, 2025. 2
- [8] Joshua Niemeijer, Jan Ehrhardt, Timo Kepp, Jörg P Schäfer, and Heinz Handels. Overcoming the sensor delta for semantic segmentation in oct images. In *Medical Imaging 2023: Computer-Aided Diagnosis*, pages 217–223. SPIE, 2023. 2
- [9] Joshua Niemeijer, Jan Ehrhardt, Hristina Uzunova, and Heinz Handels. Tsynd: Targeted synthetic data generation for enhanced medical image classification: Leveraging epistemic uncertainty to improve model performance. In *International Workshop on Simulation and Synthesis in Medical Imaging*, pages 69–78. Springer, 2024. 2
- [10] Taesung Park, Alexei A. Efros, Richard Zhang, and Jun-Yan Zhu. Contrastive learning for unpaired image-to-image translation. In *Computer Vision – ECCV 2020*, pages 319–345, Cham, 2020. Springer International Publishing. 6
- [11] Walter H. L. Pinaya, Petru-Daniel Tudosiu, Jessica Dafflon, Pedro F. Da Costa, Virginia Fernandez, Parashkev Nachev, Sebastien Ourselin, and M. Jorge Cardoso. Brain imaging generation with latent diffusion models. In *Deep Generative Models*, pages 117–126, Cham, 2022. Springer Nature Switzerland. 5
- [12] 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*, pages 10684–10695, 2022. 3, 7
- [13] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 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*, pages 234–241. Springer, 2015. 6
- [14] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. ACDC: The Adverse Conditions Dataset With Correspondences for Semantic Driving Scene Understanding. In *Proc. of ICCV*, pages 10765–10775, virtual, 2021. 5
- [15] Marc S Seibel, Joshua Niemeijer, Marc Rowedder, Helge Sudkamp, Timo Kepp, Gereon Hüttmann, and Heinz Handels. Reducing the impact of domain shift in deep learning for oct segmentation using image manipulations. In *Medical Imaging 2024: Computer-Aided Diagnosis*, pages 293–302. SPIE, 2024. 2
- [16] Marc S Seibel, Hristina Uzunova, Timo Kepp, and Heinz Handels. Anatomical conditioning for contrastive unpaired image-to-image translation of optical coherence tomography images. In *2024 IEEE International Symposium on Biomedical Imaging (ISBI)*, pages 1–5. IEEE, 2024. 2
- [17] Marc Steffen Seibel, Marc Rowedder, Julia Andresen, Richard Neffin, Tobias Neumann, Helge Sudkamp, Heinz Handels, and Timo Kepp. Enhancing retinal self-oct image quality: a deep-learning-based pipeline. In *Medical Imaging 2025: Image Processing*, pages 249–254. SPIE, 2025. 5
- [18] Nick Stracke, Stefan Andreas Baumann, Joshua M. Susskind, Miguel Angel Bautista, and Björn Ommer. Ctrloralter: Conditional loradapter for efficient 0-shot control & altering of t2i models, 2024. 2
- [19] Hristina Uzunova, Jan Ehrhardt, and Heinz Handels. Memory-efficient gan-based domain translation of high resolution 3d medical images. *Computerized Medical Imaging and Graphics*, 86:101801, 2020. 2
- [20] Hristina Uzunova, Leonie Basso, Jan Ehrhardt, and Heinz Handels. Synthesis of annotated pathological retinal oct data with pathology-induced deformations. In *Medical Imaging 2022: Computer-Aided Diagnosis*, pages 922–927. SPIE, 2022. 2
- [21] Claus von der Burchard, Helge Sudkamp, Jan Tode, Cristoph Ehlken, Konstantine Purtskhvanidze, Moritz Moltmann, Britta Heimes, Peter Koch, Michael Münst, Malte Vom Endt, et al. Self-examination low-cost full-field optical coherence tomography (self-oct) for neovascular age-related macular degeneration: a cross-sectional diagnostic accuracy study. *BMJ open*, 12(6):e055082, 2022. 2
- [22] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models.(2023). *arXiv preprint arXiv:2308.06721*, 2023. 2
- [23] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3836–3847, 2023. 2, 3

- [24] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251, 2017. [2](#), [6](#)