

META CONTROLNET: ENHANCING TASK ADAPTATION VIA META LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Diffusion-based image synthesis has attracted extensive attention recently. In particular, ControlNet that uses image-based prompts exhibits powerful capability in image tasks such as canny edge detection and generates images well aligned with these prompts. However, vanilla ControlNet generally requires extensive training of around 5000 steps to achieve a desirable control for a single task. Recent context-learning approaches have improved its adaptability, but mainly for edge-based tasks, and rely on paired examples. Thus, two important open issues are yet to be addressed to reach the full potential of ControlNet: (i) zero-shot control for certain tasks and (ii) faster adaptation for non-edge-based tasks. In this paper, we introduce a novel Meta ControlNet method, which adopts the task-agnostic meta learning technique and features a new layer freezing design. Meta ControlNet significantly reduces learning steps to attain control ability from 5000 to 1000. Further, Meta ControlNet exhibits direct zero-shot adaptability in edge-based tasks without any finetuning, and achieves control within only 100 finetuning steps in more complex non-edge tasks such as Human Pose.¹

1 INTRODUCTION

Image synthesis (Dhariwal & Nichol, 2021; Ramesh et al., 2022; Brock et al., 2018) is a rapidly growing field in computer vision and draws significant interest from various application domains. As a key approach in this area, Generative Adversarial Networks (GANs) (Brock et al., 2018; Karras et al., 2019; Goodfellow et al., 2020) employ a discriminator-generator pair, where the generator is trained to creating enhanced images via sharpening the discriminator. However, such an adversarial approach typically has difficulty to model more complex distributions. Recently, diffusion models (Dhariwal & Nichol, 2021; Ramesh et al., 2022) have emerged as a powerful alternative, excelling in high-quality image generation. These models utilize a series of denoising autoencoders to progressively refine an image from pure Gaussian noise. Among these, a new model known as Stable Diffusion (Rombach et al., 2022) has been proposed, which has better computational efficiency. Unlike traditional methods, Stable Diffusion uses latent representations for image compression, and achieves superior image quality, which includes advancements in text-to-image synthesis and unconditional image generation.

ControlNet (Zhang et al., 2023) further advances image synthesis with enhanced control over image content by using conditional control as different tasks. This approach clones the encoder and middle block of Stable Diffusion, and introduces zero convolution to link with the decoders of Stable Diffusion. Such a setup allows ControlNet to accept image prompt inputs, such as canny or HED edge, and can generate images specific to certain tasks, demonstrating improved control from both image and textual inputs. However, ControlNet’s capability for precise control requires extensive training. Specifically, learning to control a new task demands about 5000 steps. Recently, Prompt Diffusion was proposed in (Wang et al., 2023), which leverages in-context learning idea to enhance ControlNet’s adaptability to new tasks, but requires task-specific example pairs for training.

Although ControlNet and its variants have achieved enhanced the generalizability, several critical open issues remain unresolved to reach the full power of ControlNet. **Firstly**, *zero-shot* capability of ControlNet has not yet been explored, leaving it as an open question whether it is possible to control new tasks without finetuning samples. **Secondly**, while most existing studies have focused on

¹Our code is provided at https://anonymous.4open.science/r/meta_controlnet-4E0C




Figure 1: Trained from stable diffusion initial θ_{SD} , the meta learned initial θ_{meta} is used for various task adaptation.

edge-based tasks, rapid adaptation in more complex scenarios, such as the human pose task, has not yet been achieved.

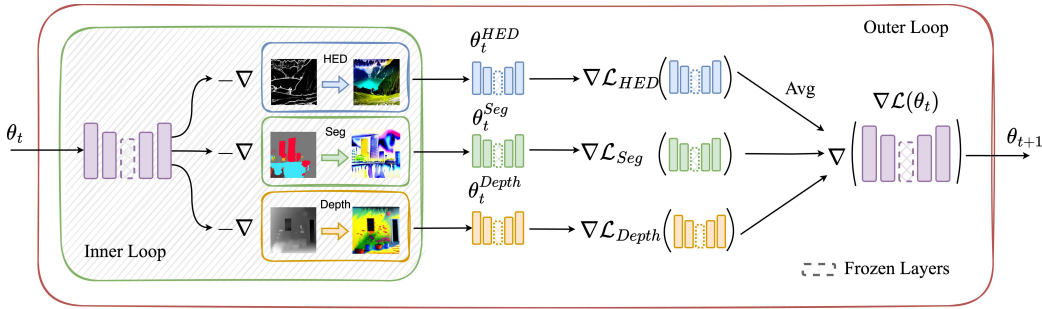


Figure 2: Meta ControlNet training pipeline. ControlNet parameter is meta updated via meta tasks (HED, Segmentation, Depth). Stable Diffusion parameters are fixed and ControlNet middle layers (Encoder Block 4 and Middle Block) are frozen during the training phase.

In this paper, we propose a novel **Meta ControlNet** method to address the aforementioned open issues. Specifically, Meta ControlNet adopts the FO-MAML (Finn et al., 2017) framework with various image condition types serving as different meta tasks. The inner-loop training of Meta ControlNet takes finetune steps separately for each task. Then the outer-loop training updates meta parameters (i.e., the model initial) based on averaged gradients over all training tasks.

(Novel Layer Freezing Design) Meta ControlNet features a new layer freezing design. Typically, meta learning algorithms such as ANIL (Raghu et al., 2019) freezes the earlier embedding layers during the inner-loop training. As a sharp difference, Meta ControlNet freezes *latter* encoder block and the *middle* block during meta training. This idea is based on the observation that the initial encoder blocks are directly linked to the control images of control tasks. It is essential to finetune these encoder blocks for individual task. On the other hand, the middle and latter encoder blocks, which capture common and high-level information, can be retained and shared across tasks. Such a design has been proven to be critical for Meta ControlNet to exhibit desirable performance in our experiments. Note that for the meta testing phase, we recommend training all layers to achieve the best possible adaptation performance.

In the following, we highlight the superior experimental performance that Meta ControlNet achieves:

- **Fast Learning of Control Ability:** Our proposed design significantly enhances the efficiency of ControlNet’s learning process. Our experiments demonstrate that Meta ControlNet acquires control abilities within only 1000 steps, a stark improvement over the vanilla ControlNet that achieves the same ability with 5000 steps. Meanwhile, this efficiency is demonstrated across three meta training tasks, showcasing the method’s versatility.
- **Zero-Shot Adaptation for Edge-based Tasks:** Meta ControlNet produces generalizable model initial, which exhibits exceptional control capabilities in zero-shot settings, especially for edge-

108 based tasks such as the canny task. This indicates that our model can adapt to new edge-based
109 tasks without any task-specific finetuning. This is the *first* achievement of successful zero-shot
110 adaptation by ControlNet. Our experiments also indicate that few-shot finetuning often further
111 enhances the fidelity of generated images.

- 112 • **Fast Adaptation in Non-edge Tasks:** For challenging tasks in few-shot contexts, our learned
113 model initial exhibits a robust adaptation ability. For instance, it can adapt to the human pose
114 task within merely 100 steps and excel in the more complex human pose mapping task in only
115 200 steps. These achievements not only surpass all existing benchmarks but also substantially
116 reduce image sample number required for adaptation.

118 2 RELATED WORK

119 2.1 DIFFUSION MODEL AND CONTROLNET

120 With the recent advancement of score-based generative models (Ho et al., 2020; Song & Ermon,
121 2019; Song et al., 2020b;a), diffusion models have achieved remarkable performance in text-to-image
122 synthesis (Dhariwal & Nichol, 2021; Ramesh et al., 2022). Essentially, the diffusion model learns a
123 time-varying mapping that gradually transforms a random noise into the sample space via a reverse
124 diffusion process (Ho et al., 2020). Stable Diffusion (SD) (Rombach et al., 2022), as an important
125 step to achieve high-resolution image generation, utilizes a variational autoencoder to first encode
126 images into a latent space, and then learns a time-conditioned U-Net to perform the denoising process
127 on the latent space based on text prompts.

130 To allow diffusion models to receive more diverse user-specific guidance for image generation,
131 Composer (Huang et al., 2023), ControlNet (Zhang et al., 2023), GLIGEN (Li et al., 2023b) and T2I-
132 Adapter (Mou et al., 2023) were proposed as general approaches to introduce additional controlling
133 signals. Among these, ControlNet (Zhang et al., 2023) stands out by its superior performance in
134 various downstream tasks ranging from sketch (edge, skeleton) to geometry (depth, normal) guided
135 generation. Technically, ControlNet freezes the original SD model, while finetuning a duplicate of the
136 pre-trained SD to integrate additional control signals via zero-initialized convolution modules. Such
137 an adaptation scheme significantly reduces training costs by re-using the image prior learned in the
138 pre-trained SD. ControlNet-XS (Rother, 2023) further investigates the size and architectural design
139 of ControlNet and proposes a more parameter-efficient architecture. As downstream applications, by
140 leveraging ControlNet, Goel et al. (2023) is able to edit the structure and appearance properties for
141 each object in the image, and Chu et al. (2023); Wu et al. (2023) enforces temporal consistency for
142 video generation, Ma et al. (2023); Seo et al. (2023) make pre-trained SD aware of 3D knowledge and
143 multi-view geometry. However, for each of these tasks, an independent adapter is required for each
144 condition. The modified version Multi-ControlNet (Zhang et al., 2023) demonstrates the possibility of
145 composing multiple tasks. Uni-ControlNet (Zhao et al., 2023) proposes a unified framework allowing
146 for the simultaneous utilization of different local and global controls. Prompt Diffusion (Wang et al.,
147 2023) trains an open-domain ControlNet in an in-context learning manner. However, none of these
148 studies are designed for or have been demonstrated to have zero-shot generalization capability for
149 unseen control tasks.

150 2.2 META LEARNING

151 Meta learning focuses on few-shot learning scenarios, aiming to develop algorithms that leverage a
152 large set of pre-defined tasks to improve performance on unseen instances with only a few or even
153 zero extra training data samples. In this paper, we focus on “learning-to-initialize” approaches such as
154 Model-Agnostic Meta Learning (MAML) (Finn et al., 2017), which learns an initialization point from
155 which models can fast adapt to new tasks. MAML (Finn et al., 2017) algorithm involves two layers of
156 training, where at each iteration, the inner loop optimizes the parameter for each task independently
157 starting from the current initialization, whereas the outer loop estimates the gradient with respect
158 to the inner loop optimization path to update the initialization. Since differentiating through an
159 optimization algorithm is often computationally expensive, FO-MAML (Finn et al., 2017) proposes to
160 simplify the outer-loop gradient computation by directly averaging task-specific gradients evaluated
161 at the outputs of the inner loop. ANIL (Raghu et al., 2019) improves MAML via feature reusing,
where the main feature backbone is frozen and only prediction heads are updated by each task in the

inner loop. Reptile (Nichol et al., 2018; Nichol & Schulman, 2018) simplifies the process by aiming for an initialization that minimizes the expected loss across all tasks, similar to joint training.

In computer vision, a popular application of meta learning is few-shot image classification, where a network is adapted to new classes using only a small number of labeled instances per class (Finn et al., 2017; Ravi & Larochelle, 2016; Verma et al., 2020). Li et al. (2018) also adopts MAML-based methods to improve cross-domain generalization. MetaGAN (Zhang et al., 2018) and MetaDiff (Zhang & Yu, 2023) combines MAML respectively with GAN (Goodfellow et al., 2020) and diffusion model (Song & Ermon, 2019; Ho et al., 2020) to facilitate few-shot image classification. More recently, meta learning has been employed to accelerate implicit neural representation for visual signals (Sitzmann et al., 2020; Tancik et al., 2021).

3 META CONTROLNET

In this section, we first propose the Meta ControlNet method, and then explain how to select and structure both training and adaptation tasks.

3.1 ALGORITHM DESIGN

In this section, we propose our algorithm Meta ControlNet, which maintains the Stable Diffusion network while training its duplicates via the task-agnostic meta learning technique for obtaining adaptive model initial.

In particular, Meta ControlNet adopts three control tasks (HED, Segmentation, Depth) as the primary meta tasks. The training of Meta ControlNet takes the double-loop training framework of FO-MAML (Finn et al., 2017), as depicted in Figure 2, and is described in detail as follows.

The *inner-loop* training of Meta ControlNet takes finetune steps separately for each task. During each step t , the meta parameter θ_t of Meta ControlNet is finetuned independently for each task based on gradient descent as follows:

$$\text{(Inner Loop)} \quad \theta_t^{task} = \theta_t - \alpha \nabla \mathcal{L}_{task}(\theta_t),$$

where $task \in \{\text{HED, Seg, Depth}\}$ and α represents the step size. Note that we here update the parameter only once in the inner loop to enhance efficiency.

The *outer-loop* training first calculates the meta gradient $\nabla \mathcal{L}(\theta_t)$ as follows by taking an average of the gradients across all tasks, based on each task’s finetuned parameters in the inner loop:

$$\nabla \mathcal{L}(\theta_t) = \text{Avg}_{task}(\nabla \mathcal{L}_{task}(\theta_t^{task})),$$

where "Avg" denotes the averaging operator over all task gradients. Then the meta parameter θ_t is updated using the meta gradient as follows:

$$\text{(Outer Loop)} \quad \theta_{t+1} = \theta_t - \alpha \nabla \mathcal{L}(\theta_t),$$

where α is the step size. This design guides Meta ControlNet to minimize the loss of the finetuned model for each task, and hence makes the model more responsive to updates and enables fast adaptability.

Novel Layer Freezing Design: A main novel component that Meta ControlNet features is the design of freezing layers during training, which turns out to be critical for its superior performance. Typically, meta learning algorithms such as ANIL (Raghu et al., 2019) freeze the earlier embedding layers during the inner-loop training. Meta ControlNet has sharp differences in two aspects. **Firstly**, Meta ControlNet freezes *latter* encoder block and the *middle* block during meta training. This idea is based on the observation that the initial encoder blocks are directly linked to the control images of control tasks. Given that our Stable Diffusion initial was used to process Gaussian noise rather than control task-specific inputs, and different control image styles signify distinct tasks, it is essential to finetune these initial encoder blocks for individual tasks. On the other hand, the middle and latter encoder blocks, which capture common and high-level information, can be retained and shared across tasks. Therefore, during the training process, Meta ControlNet selectively freezes the Encoder Block4 and Middle block of the U-Net, and focuses on training the remaining parameters. The detailed architecture is available in Appendix. **Secondly**, unlike ANIL, where freezing occurs only in the



Figure 3: Validation set samples from each training task (HED, Depth, Segmentation) after 1000 steps of training updates.

inner loop, Meta ControlNet applies layer freezing in both inner and outer loops to achieve better efficiency by leveraging our network’s initial high-quality image generation capability from Stable Diffusion.

It is important to note that the meta design and layer freezing are applied only during Meta ControlNet’s training phase. During the adaptation phase, finetuning uses the standard ControlNet training protocol without using layer freezing or meta learning methods.

3.2 TASK SELECTION

Training Tasks: During training phase, we choose HED, Segmentation, and Depth map control as our training tasks. Specifically, we obtain HED map using HED boundary detector proposed by (Xie & Tu, 2015). We obtain Segmentation map using Uniformer (Li et al., 2023a). We collect Depth map using Midas (Ranftl et al., 2020).

Adaptation Tasks: Following Wang et al. (2023), we utilize Canny Edge maps and Normal maps as our adaptation tasks. These tasks, which align the generated image with the control image’s edges, are categorized as *edge-based* tasks. Additionally, we introduce two more complex tasks to demonstrate the versatility of our model: Human Pose (line segments to objects) and Human Pose Mapping (objects to line segments), referred to as *non-edge* tasks.

In detail, we collect Canny Edge by using Canny Edge detector (Canny, 1986). We obtain Normal map by applying Midas (Ranftl et al., 2020). We collect human pose and its reverse human pose mapping by using Openpose (Cao et al., 2019).

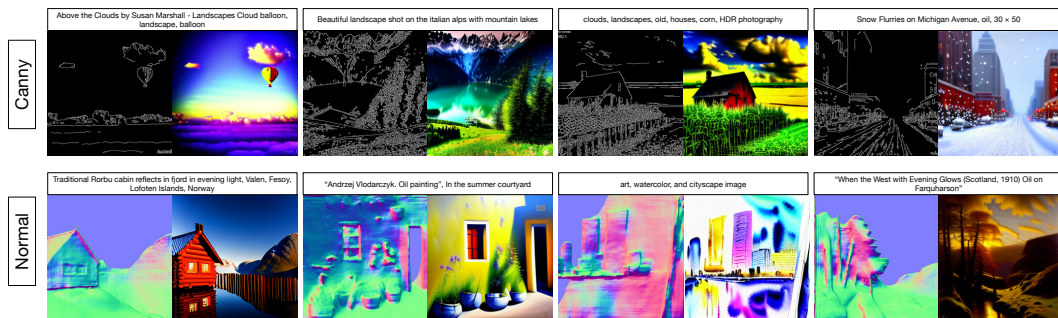

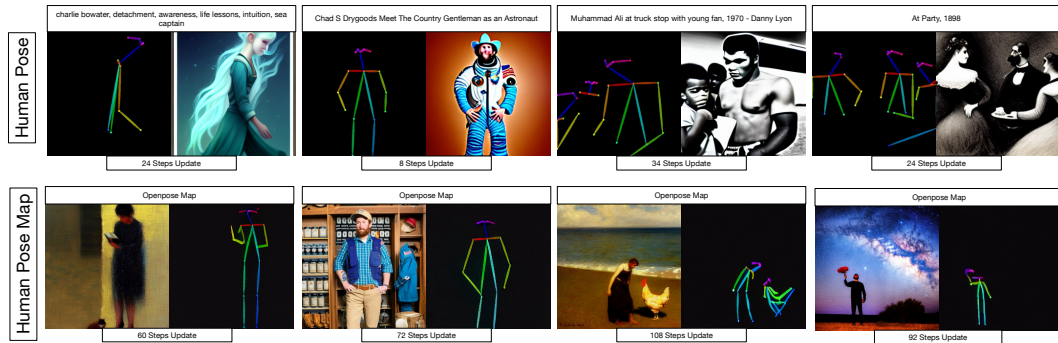


Figure 4: Samples from edge-based tasks (Canny, Normal) in zero-shot adaptation.



282 Figure 5: Sample comparison between proposed Meta ControlNet and Prompt Diffusion (PD) baseline
283 for canny task in few-shot finetuning. PD requires example pairs to update and thus is only available
284 in odd number few-shot setting.



298 Figure 6: Samples from the validation set for non-edge tasks (Human Pose, Human Pose Map) in a
299 finetuning context. Below each image, the number of updates indicates the first iteration achieving
300 significant control. The validation set is evaluated every two updates.

302 4 EXPERIMENTAL RESULTS

304 **Dataset:** We use the generated CLIP-filtered dataset proposed by InstructPix2Pix (Brooks et al.,
305 2023) as our training and validation datasets. The CLIP-filtered dataset contains 313k image-prompt
306 pairs.

307 **Implementations:** Meta ControlNet is developed using the ControlNet codebase (Zhang et al., 2023),
308 and utilizes the Stable Diffusion v1.5 checkpoint for finetuning. We fix learning rate to be 1×10^{-4}
309 and batch size to be 256, and accumulate gradients over every 4 batches. Our model is trained on 4
310 Nvidia A100 GPUs. In our study, both our Meta ControlNet and the baseline of Prompt Diffusion
311 (PD) (Wang et al., 2023) are evaluated at the 8000-step checkpoint. Note that more finetuning steps
312 enhance image quality.

313 Regarding the meta design, the meta training images in the inner loop are reused in the meta
314 testing phase in the outer loop for each task in order to optimize memory efficiency. Note that the
315 aforementioned batch size of 256 is the total number of images from all tasks. Namely, we randomly
316 sample 256 images across all tasks, and organize them into batches respectively for each task. This
317 strategy ensures that our algorithm does not require additional image sample during the training
318 phase.

320 4.1 FAST CONTROL ACQUIRING IN TRAINING

321
322 The proposed Meta ControlNet is trained on tasks of HED, segmentation, and depth mapping.
323 Figure 3 displays the validation results after 1000 steps. Clearly, our Meta ControlNet generates the
images that closely match the control images with high fidelity. While the vanilla ControlNet requires

5000 steps to exhibit control ability on a single task, our algorithm Meta ControlNet enjoys rapid learning within only 1000 steps. In fact, for most images, control ability of Meta ControlNet occurs within 500 steps or even fewer, with additional training serving to enhance image fidelity.

4.2 ZERO-SHOT CAPABILITY FOR EDGE-BASED TASKS

We evaluate Meta ControlNet on edge-based tasks, specifically Canny and Normal tasks. Figure 4 presents the control images alongside the corresponding generated images and text prompts. This adaptation is assessed in a *zero-shot* context, and does not require finetuning or additional data. This is the *first* ControlNet-type method featuring *zero-shot* capability. In contrast, the baseline Prompt Diffusion (PD) method (Wang et al., 2023) relies on example pairs for learning, rendering it unsuitable for zero-shot settings. The zero-shot results clearly showcase the superior adaptation capability of Meta ControlNet with strong control ability and high fidelity in both tasks.

Further, we compare our Meta ControlNet and the PD baseline with both finetuned in a few-shot context, namely, each method is updated with an equal number of few-shot images. Note that the proposed Meta ControlNet needs only one sample per update step, while PD requires two examples per step. The results in Figure 5 indicate that image quality by Meta ControlNet is enhanced with additional shots. Further, our Meta ControlNet clearly outperforms PD, although the fidelity of the PD gradually improves with more shots. When comparing both methods over the same number of finetuning steps, such as 10 shots for our proposed method and 21 shots for the PD baseline, our approach consistently yields higher quality images. Note that PD requires example pairs to update and thus is only available in odd number few-shot settings.

We highlight that in our experiments, most images generated in zero-shot already exhibit high fidelity, and thus additional few-shot finetuning is not required and might not result in substantial improvements.


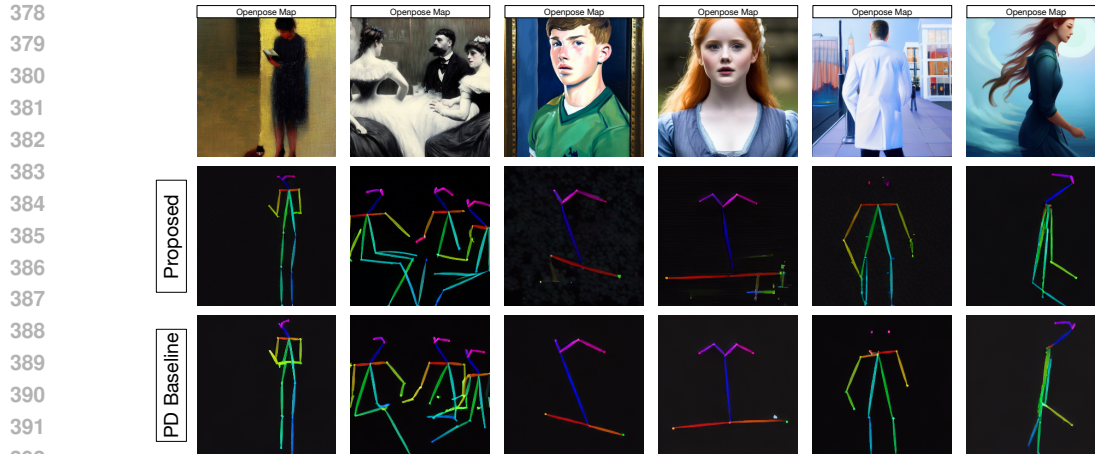


Figure 7: Validation sample comparison between proposed Meta ControlNet and Prompt Diffusion (PD) baseline for Human Pose task after 100 steps of finetuning updates.

4.3 FAST ADAPTATION FOR NON-EDGE TASKS

We evaluate the generalizability of Meta ControlNet in non-edge tasks, with focus on human pose and its reverse mapping. These non-edge tasks, which are typically challenging in few-shot setting, require a training-like approach for finetuning. To enhance stability, we double the gradient accumulation from 4 to 8 and keep the batch size to be 256.

For Meta ControlNet, we assess validation at every two steps, and record the first instance when the generated image is aligned with the control image, as shown in Figure 6. We observe that for the human pose task, effective control is achieved within 50 steps. The human pose mapping task,



393 Figure 8: Validation sample comparison between proposed Meta ControlNet and Prompt Diffusion
394 (PD) baseline for Human Pose Mapping task after 200 steps of finetuning updates.
395



411 Figure 9: Sample comparison among different freeze methods for edge-based tasks in zero-shot
412 context. En_N refers to the N^{th} Encoder Block. 'Freeze $En_4 + \text{Middle}$ ' refers to freezing the 4th
413 Encoder Block and the Middle Block in U-Net, which is adapted in Meta ControlNet.
414

415
416 which converts human poses into line segments, presents a greater challenge due to the deviation
417 from the high-quality images typically generated by stable diffusion. Nevertheless, Meta ControlNet
418 demonstrates control over most samples within approximately 100 steps. We note that in tasks such
419 as human pose mapping, minor errors can occur, for example, incorrectly depicting a chicken pose
420 in the third sample pair. This is due to ControlNet's inherent limitation in distinguishing between
421 human and animal, a distinction that requires learning from more samples.

422 To compare our method with the PD baseline in the human pose task, we evaluate 100-step finetuning
423 in Figure 7. Despite the use of example pairs, PD achieves control but at the cost of reduced fidelity.
424 In contrast, Meta ControlNet maintains both high control and fidelity. For the more challenging
425 human pose mapping task, our method achieves comparable results as the PD baseline with the same
426 200 steps, but with only half number of images used by PD, demonstrating our better efficiency.

427 We note that the different convergence speeds to reach control between the human pose task and
428 its mapping counterpart, i.e., 100 versus 200 steps, arise from the inherent characteristics of Stable
429 Diffusion and our choice of training tasks. Both the standard Stable Diffusion and our selected
430 tasks are geared towards generating natural images rather than line segments. Consequently, in the
431 adaptation phase, the model more readily adapts to the human pose task, which involves creating
natural human images, as opposed to learning to generate line segments from human images, a

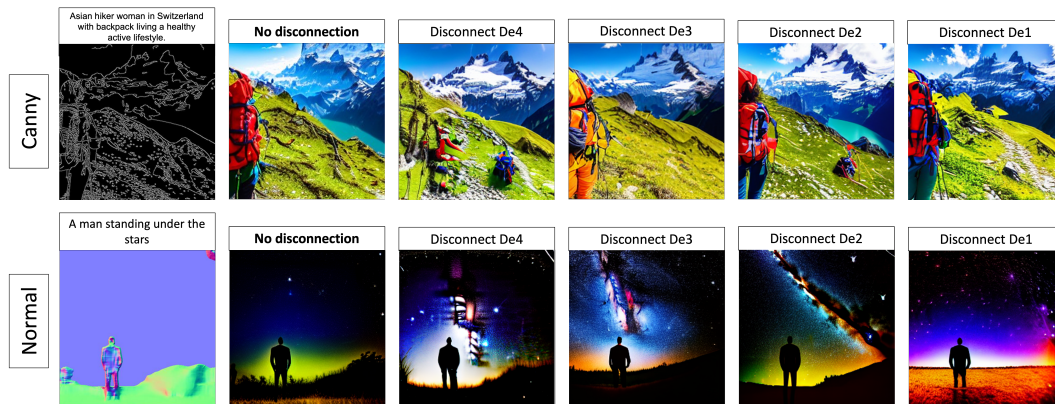


Figure 10: Sample comparison among various connection methods for zero-shot edge-based tasks. De_N refers to the N^{th} Decoder Block.

requirement of the mapping task. Nevertheless, our method still achieves better efficiency than PD, namely, comparable results but with only half number of images during training.

5 ABLATION STUDY

5.1 LAYER FREEZING

To evaluate the effectiveness of various freezing designs, we conduct an experimental comparison focusing on the canny and normal tasks in a zero-shot setting. The experiment compares our freezing strategy (freezing Encoder Block 4 and the Middle block in U-Net) against other methods: freezing Encoder Blocks 2 to 4 plus the Middle block, freezing Encoder Blocks 1 to 3, and no freezing. The U-Net architecture is available in Appendix. The comprehensive results are detailed in Figure 9, with all methods under identical experimental conditions.

Our results indicate that while all freezing designs facilitate control ability, freezing Encoder Block 4 and the Middle block achieves the most accurate alignment with the control image and the highest image fidelity. This superiority is likely because the first three encoder blocks in U-Net are more task-specific, as they are more directly connected to the control image. In contrast, the latter encoder block and the Middle block contribute significantly to the high-quality image output inherent in stable diffusion, making them ideal candidates for freezing across diverse tasks.

5.2 DECODER CONNECTION

We evaluate our algorithm using different types of connections between the ControlNet decoder and the pre-trained Stable Diffusion (SD) model decoder. Specifically, we use Meta ControlNet with all decoder connected by zero convolution as our baseline. We then disconnect Decoder Blocks 4 to 1 from the pre-trained SD model in separate variants (each Decoder Block n corresponds to Encoder Block n). These variants are trained under the same conditions as the baseline, up to the 8000-step checkpoint, and are then tested on normal and canny edge tasks. Figure 10 depicts the result and suggests that disconnecting Decoder Block 4 significantly reduces image fidelity, and often results in images with repetitive lines or objects with strange shapes. In contrast, disconnecting the other three blocks produces results similar to our baseline. This indicates that Decoder Block 4 plays a critical role in ensuring high image fidelity, which is consistent with our design choice of freezing Encoder Block 4 for this purpose.

6 CONCLUSION

In our study, we propose a novel Meta ControlNet approach, which adopts the meta learning technique and features novel freezing layer design to learn a generalizable ControlNet initial. This method

486 exhibits rapid training convergence, and requires only 1000 steps to effectively control generative
487 imaging. Further, such a meta initial exhibits remarkable zero-shot adaptability for edge-based tasks,
488 the first demonstration in this domain. It also excels in more challenging non-edge tasks, and adapts
489 rapidly within 100 steps for the human pose task and 200 steps for the human pose map task. These
490 results not only outperform existing baselines in terms of control ability and efficiency but also
491 represent significant advancement beyond vanilla ControlNet.

492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

REFERENCES

- 540
541 Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural
542 image synthesis. In *International Conference on Learning Representations (ICLR)*, 2018.
543
- 544 Tim Brooks, Aleksander Holynski, and Alexei A. Efros. Instructpix2pix: Learning to follow image
545 editing instructions, 2023.
546
- 547 John Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis
548 and Machine Intelligence*, PAMI-8(6):679–698, 1986. doi: 10.1109/TPAMI.1986.4767851.
549
- 550 Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: Realtime
551 multi-person 2d pose estimation using part affinity fields, 2019.
- 552 Ernie Chu, Shuo-Yen Lin, and Jun-Cheng Chen. Video controlnet: Towards temporally consistent
553 synthetic-to-real video translation using conditional image diffusion models. *arXiv preprint
554 arXiv:2305.19193*, 2023.
- 555 Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances
556 in neural information processing systems*, 34:8780–8794, 2021.
557
- 558 Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of
559 deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.
560
- 561 Vidit Goel, Elia Peruzzo, Yifan Jiang, Dejia Xu, Nicu Sebe, Trevor Darrell, Zhangyang Wang, and
562 Humphrey Shi. Pair-diffusion: Object-level image editing with structure-and-appearance paired
563 diffusion models. *arXiv preprint arXiv:2303.17546*, 2023.
- 564 Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair,
565 Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the
566 ACM*, 63(11):139–144, 2020.
567
- 568 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
569 neural information processing systems*, 33:6840–6851, 2020.
- 570 Lianghua Huang, Di Chen, Yu Liu, Yujun Shen, Deli Zhao, and Jingren Zhou. Composer: Creative
571 and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*,
572 2023.
- 573 Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative
574 adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
575 recognition (CVPR)*, pp. 4401–4410, 2019.
576
- 577 Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning
578 for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*,
579 volume 32, 2018.
- 580 Kunchang Li, Yali Wang, Junhao Zhang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and
581 Yu Qiao. Uniformer: Unifying convolution and self-attention for visual recognition, 2023a.
582
- 583 Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li,
584 and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the
585 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22511–22521, 2023b.
586
- 587 Wufei Ma, Qihao Liu, Jiahao Wang, Angtian Wang, Yaoyao Liu, Adam Kortylewski, and Alan Yuille.
588 Adding 3d geometry control to diffusion models. *arXiv preprint arXiv:2306.08103*, 2023.
- 589 Chong Mou, Xintao Wang, Liangbin Xie, Jian Zhang, Zhongang Qi, Ying Shan, and Xiaoju Qie.
590 T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion
591 models. *arXiv preprint arXiv:2302.08453*, 2023.
592
- 593 Alex Nichol and John Schulman. Reptile: a scalable metalearning algorithm. *arXiv preprint
arXiv:1803.02999*, 2(3):4, 2018.

- 594 Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv*
595 *preprint arXiv:1803.02999*, 2018.
- 596
- 597 Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. Rapid learning or feature reuse?
598 towards understanding the effectiveness of maml. *arXiv preprint arXiv:1909.09157*, 2019.
- 599
- 600 Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-
601 conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- 602 RenÅl Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards
603 robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer, 2020.
- 604
- 605 Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *International*
606 *conference on learning representations*, 2016.
- 607
- 608 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-
609 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF confer-*
ence on computer vision and pattern recognition, pp. 10684–10695, 2022.
- 610
- 611 Denis Zavadski Carsten Rother. Controlnet-xs. <https://vislearn.github.io/ControlNet-XS/>, 2023.
- 612
- 613 Junyoung Seo, Wooseok Jang, Min-Seop Kwak, Jaehoon Ko, Hyeonsu Kim, Junho Kim, Jin-Hwa
614 Kim, Jiyoung Lee, and Seungryong Kim. Let 2d diffusion model know 3d-consistency for robust
text-to-3d generation. *arXiv preprint arXiv:2303.07937*, 2023.
- 615
- 616 Vincent Sitzmann, Eric Chan, Richard Tucker, Noah Snaveley, and Gordon Wetzstein. Metasdf:
617 Meta-learning signed distance functions. *Advances in Neural Information Processing Systems*, 33:
10136–10147, 2020.
- 618
- 619 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*
620 *preprint arXiv:2010.02502*, 2020a.
- 621
- 622 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
623 *Advances in neural information processing systems*, 32, 2019.
- 624
- 625 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
626 Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint*
arXiv:2011.13456, 2020b.
- 627
- 628 Matthew Tancik, Ben Mildenhall, Terrance Wang, Divi Schmidt, Pratul P Srinivasan, Jonathan T
629 Barron, and Ren Ng. Learned initializations for optimizing coordinate-based neural representations.
630 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
2846–2855, 2021.
- 631
- 632 Vinay Kumar Verma, Dhanajit Brahma, and Piyush Rai. Meta-learning for generalized zero-shot
633 learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 6062–
6069, 2020.
- 634
- 635 Zhendong Wang, Yifan Jiang, Yadong Lu, Yelong Shen, Pengcheng He, Weizhu Chen, Zhangyang
636 Wang, and Mingyuan Zhou. In-context learning unlocked for diffusion models, 2023.
- 637
- 638 Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu,
639 Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion
640 models for text-to-video generation. In *Proceedings of the IEEE/CVF International Conference on*
Computer Vision, pp. 7623–7633, 2023.
- 641
- 642 Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proceedings of the IEEE*
643 *international conference on computer vision*, pp. 1395–1403, 2015.
- 644
- 645 Baoquan Zhang and Demin Yu. Metadiff: Meta-learning with conditional diffusion for few-shot
646 learning. *arXiv preprint arXiv:2307.16424*, 2023.
- 647
- 648 Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image
diffusion models, 2023.

648 Ruixiang Zhang, Tong Che, Zoubin Ghahramani, Yoshua Bengio, and Yangqiu Song. Metagan: An
649 adversarial approach to few-shot learning. *Advances in neural information processing systems*, 31,
650 2018.
651
652 Shihao Zhao, Dongdong Chen, Yen-Chun Chen, Jianmin Bao, Shaozhe Hao, Lu Yuan, and Kwan-
653 Yee K Wong. Uni-controlnet: All-in-one control to text-to-image diffusion models. *arXiv preprint*
654 *arXiv:2305.16322*, 2023.
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701

A APPENDIX

A.1 DETAILED META CONTROLNET ARCHITECTURE

In Figure 11, the detailed Meta ControlNet architecture is illustrated. This architecture shows that during both training and testing phases, the stable diffusion part remain locked, in line with the ControlNet settings. Specifically, for the ControlNet part, SD Encoder Block4 and SD Middle Block are frozen during the meta training phase, while other blocks are subject to finetuning. This approach is chosen because SD Encoder Blocks 1-3 are more closely linked to the control image, necessitating their adaptation to capture task-specific differences. During the meta testing phase, all blocks within the ControlNet part are finetuned to optimize performance.

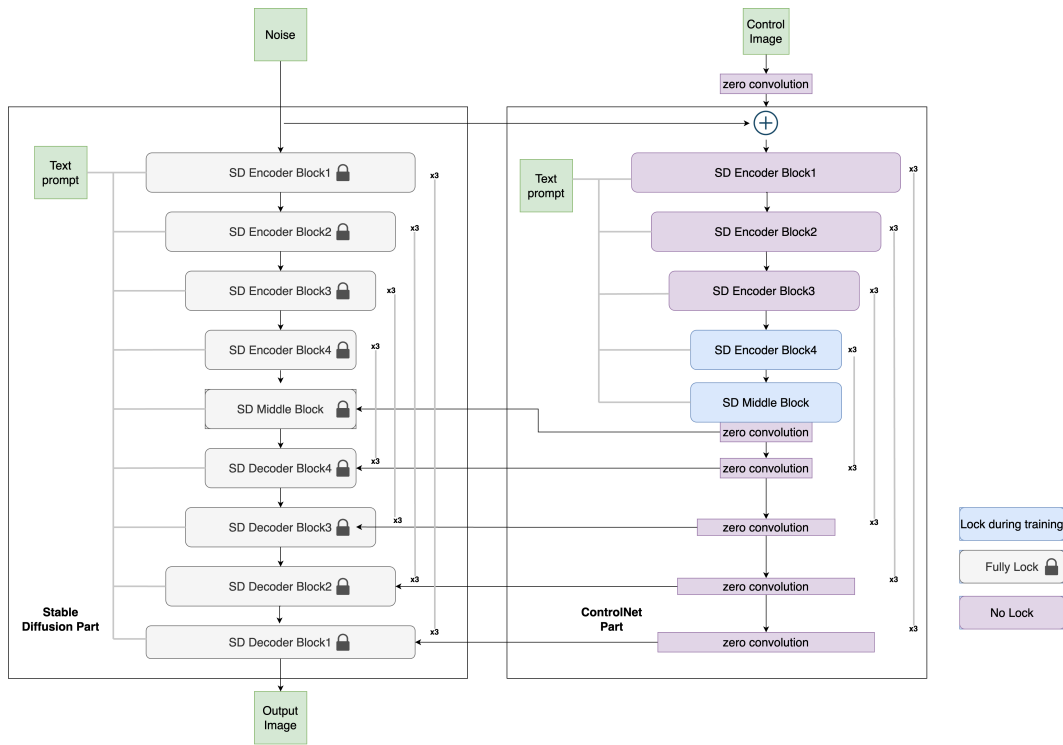


Figure 11: Detailed architecture of Meta ControlNet