

---

# CommonCanvas: An Open Diffusion Model Trained with Creative-Commons Images

---

Aaron Gokaslan<sup>1</sup> A. Feder Cooper<sup>1</sup> Jasmine Collins<sup>2</sup> Landan Seguin<sup>2</sup>

Austin Jacobson<sup>2</sup> Mihir Patel<sup>2</sup> Jonathan Frankle<sup>2</sup> Cory Stephenson<sup>2</sup> Volodymyr Kuleshov<sup>1</sup>

<sup>1</sup>Cornell Tech

{akg87,afc78,vk379}@cornell.edu

<sup>2</sup>Databricks Mosaic

{firstname.lastname}@databricks.com

## Abstract

We assemble a dataset of Creative-Commons-licensed (CC) images, which we use to train a set of open diffusion models that are qualitatively competitive with Stable Diffusion 2 (SD2). This task presents two challenges: (1) high-resolution CC images lack the captions necessary to train text-to-image generative models; (2) CC images are relatively scarce. In turn, to address these challenges, we use an intuitive transfer learning technique to produce a set of high-quality synthetic captions paired with curated CC images. We then develop a data- and compute-efficient training recipe that requires as little as 3% of the LAION data (i.e., roughly 70 million examples) needed to train existing SD2 models, but obtains the same quality. These results indicate that we have a sufficient number of CC images (also roughly 70 million) for training high-quality models. Our training recipe also implements a variety of optimizations that achieve  $\sim 3X$  training speed-ups, and that enable rapid model iteration. We leverage this recipe to train several high-quality text-to-image models, which we dub the *CommonCanvas* family. Our largest model achieves comparable performance to SD2 on human evaluation, even though we only use a CC dataset that is  $<3\%$  the size of LAION and synthetic captions for training. We release our models, data, and code at <https://github.com/mosaicml/diffusion/blob/main/assets/common-canvas.md>.

## 1 Introduction

Current methods train high-quality, text-to-image (T2I) models with. A lack of curated datasets that are large enough for the task has led researchers to turn to web-scraped solutions [29, 30], like LAION-2B [26]. The use of web-scraped data is a very common practice for training generative models, however, US courts have yet to definitively rule if this is permissible under copyright law [1, 13, 15, 20, 21, 60]. In response, recent work has begun to investigate alternative methods of navigating copyright concerns in text generation [39], code completion [16, 51], and image generation [24]. Nevertheless, matching the performance of state-of-the-art models remains a challenge. In this work, we study the following natural question: *Is it possible to efficiently produce a high-quality T2I model by training only on Creative-Commons-licensed data?*

We suggest a possible path forward, training a suite of T2I architectures using *only* open-licensed, Creative-Commons (CC) images (Figures 1 & 2). This task brings to light two significant challenges. The first problem is data incompleteness: almost all CC images lack the captions necessary to train a high-quality T2I model. The second is data scarcity: there are relatively few high-resolution CC images — roughly 70 million, compared to LAION-2B’s roughly 2 billion [26].




Figure 1: Selection of text prompts. Using entirely Creative-Commons images and our synthetic captioning approach, we achieve comparable qualitative performance to Stable Diffusion 2 (SD2-base), as seen in CommonCanvas generations, while only requiring a small fraction ( $< 3\%$ ) of the amount of training data. We include results for two CommonCanvas architectures, small (S) and large (L) (Section 6), and two CC-image datasets, commercial (C) and non-commercial (NC) (Section 4). We label our results accordingly as CommonCanvas- $\langle$ architecture $\rangle$ - $\langle$ dataset $\rangle$ .

We address the data incompleteness problem by using a pre-trained BLIP-2 model [34], which we use to produce high-quality, synthetic captions for a set of curated, open licensed CC images. This is an intuitive transfer-learning solution: leveraging powerful pre-trained generative models to produce synthetic labels for an unlabeled dataset, which we can then use to train a different multimodal generative model. We note that this is an increasingly common pattern in the literature, which we shorthand with the name *telephoning*.

To deal with data scarcity, we propose a data- and compute-efficient training recipe that obtains the same quality as SD2, but (perhaps surprisingly) requires as little as 3% of the LAION-2B data (i.e., roughly 70 million examples) originally used to train SD2. We call this model SD2-base. These results indicate that we have a sufficient number of CC images (also roughly 70 million) for training high-quality models. Our training recipe also implements a variety of optimizations that achieve  $\sim 3X$  training speed-ups, and that allow for rapid model iteration.

The above methods enable us to create *CommonCanvas*, a suite of latent diffusion model (LDM) architectures trained on our curated dataset of CC images and synthetic captions, which we denote *CommonCatalog*. For CommonCanvasL-NC, we swap SD2’s UNet for SDXL to demonstrate how even with less data, larger models do not overfit to this smaller dataset. Our largest model achieves performance comparable to SD2-base on human evaluation of Parti Prompts [66], even though our CommonCatalog training dataset is  $< 3\%$  the size of LAION and has synthetically generated captions. Figure 1 shows select samples from our CommonCanvas models compared to corresponding samples from SD2-base. Although this model is a larger and - likely - more capable model architecture than SD2, we find it surprising and important that it is possible to train an SD2-quality model at all based on such a limited dataset that was cobbled together in this fashion. This reveals a promising path forward for future research on highly-capable, open T2I models. In summary, we:

- Synthesize a set of high-quality captions for uncaptioned CC images, which we can then use together for training. We note that this type of transfer-learning technique is increasingly common, and we give it the shorthand name *telephoning* (Section 3).
- Curate *CommonCatalog*, a dataset of roughly 70 million open-licensed CC images, for which we use telephoning to generate accompanying high-quality synthetic captions (Section 4).

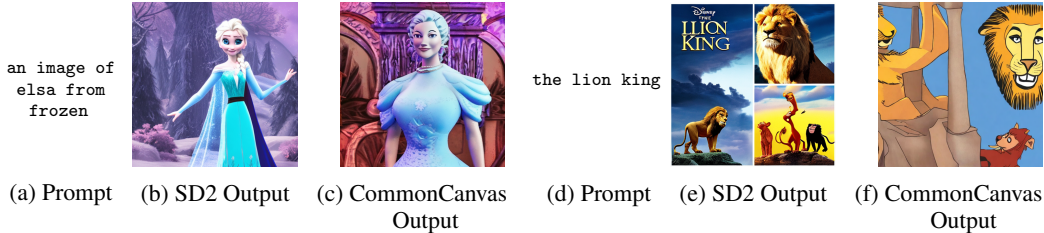


Figure 2: When given prompts for concepts related to Disney movies (a, d), SD2-base generates a recognizable image of Elsa from *Frozen* (b) and a poster-like image with a misshapen Disney logo and characters resembling those from *The Lion King* (e), and CommonCanvas (-SC) does not (c, f).

- Train and evaluate *CommonCanvas*, a suite of LDM architectures trained on CommonCatalog. We demonstrate that these models produce competitive qualitative and quantitative results compared to the SD2-base baseline (Section 6). To make this analysis tractable, we implement a variety of training optimizations, which achieve  $\sim 3X$  speed-ups in training SD2-base (Section 5).
- Release our CommonCatalog dataset of CC images and synthetic captions along with our trained CommonCanvas model at <https://github.com/mosaicml/diffusion/blob/main/assets/common-canvas.md>.

## 2 Preliminaries and Motivation

In this section, we present background on training the T2I Stable Diffusion model, originally trained on the web-scraped LAION-2B dataset. We then discuss copyright and reproducibility with respect to LAION datasets. This discussion motivates the creation of an alternative dataset composed of open licensed, CC images with synthetic captions, which we introduce in Section 4.

### 2.1 Text-to-image generative models

Text-to-image (T2I) generative models refer to large neural networks trained on paired image-caption data examples. One such family of T2I models is Stable Diffusion (SD) [47]. SD is a latent diffusion model (LDM) that converts images to latent representations and back again using Variational Autoencoders (VAEs) [23]; it uses an iterative sampling procedure [57] and trains an underlying UNet [48]. The architecture also includes a text encoder, such as the Contrastive Language-Image Pre-training (CLIP) model [43] – either the original CLIP from OpenAI [45] or its open-source counterpart, OpenCLIP [10, 18].

Stable Diffusion 2 (SD2)’s UNet has approximately 865 million trainable parameters; Stable Diffusion XL (SDXL) is larger, with 2.6 billion parameters, and has other advancements involving aspect ratio bucketing, micro-conditioning, and multiple text encoders and tokenizers. In terms of training data, the SD-family of models and OpenCLIP are both trained on subsets of the LAION-5B dataset [3, 53]. The exact training dataset for CLIP is unknown, but it is likely webscraped data [45]

### 2.2 Copyright and reproducibility in relation to LAION datasets

LAION-5B is a dataset derived from a snapshot of the Common Crawl, a massive corpus of data scraped from the web. From this snapshot, the LAION organization curated pairs of image URLs and their corresponding alt-text captions for the intended use of training T2I and image-to-text (I2T) generative models [3, 53]. In practice, T2I models are typically trained on filtered subsets of the full LAION-5B dataset (e.g. LAION-2B [26]). Training T2I models on this dataset requires visiting the URLs and downloading the associated images. There are two elements of LAION datasets that are relevant to our work:

**Copyright.** The images associated with LAION datasets have unclear *provenance*: it is often not known what the original image sources are [29, 30]. Courts have not yet decided if training on these datasets is “fair use” — an important exception in copyright [29, 33, 50, 56]. In the interim, there are several copyright lawsuits for the alleged use of LAION-5B subsets to train generative models [1, 15, 20, 61].

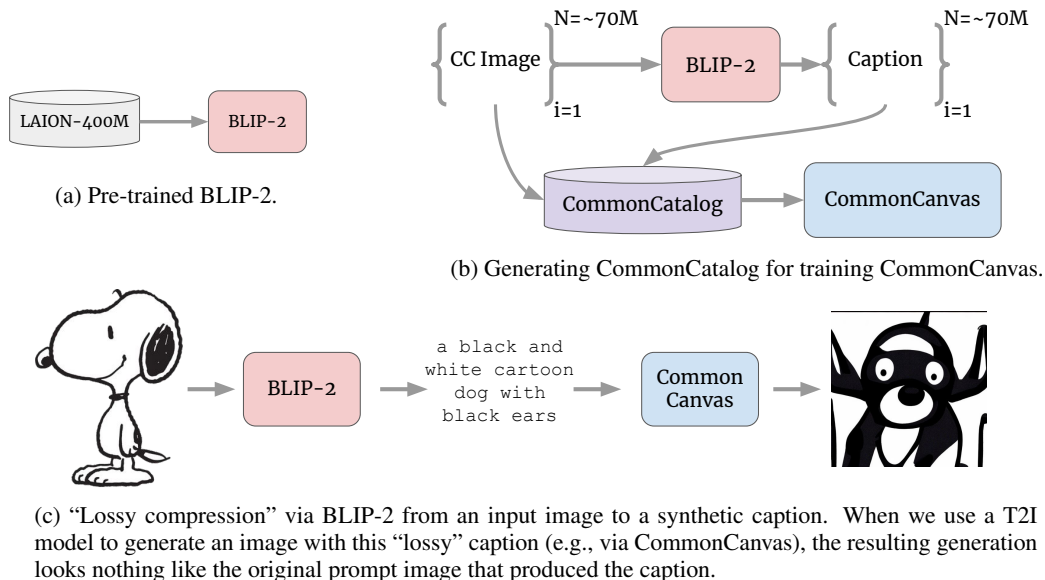


Figure 3: (a) LAION’s massive dataset of image-caption pairs is used to train BLIP-2, an image-to-text model. (b) We leverage BLIP-2 to produce synthetic captions for our caption-less CC images, and use the resulting synthetic image-caption pairs (the *CommonCatalog* dataset) to train our open diffusion model, *CommonCanvas*. (c) Although BLIP-2 was trained on LAION (e.g., including pictures of characters Snoopy), the captions it produces behave like a “lossy compression” (e.g., a black and white cartoon dog with black ears, which has no mention of Snoopy). When we supply such “lossy” captions to a T2I model, like a game of telephone, it produces outputs that no longer resemble the original images (e.g., we show how *CommonCanvas* produces an image that matches the caption, but does not look like Snoopy).

**Reproducibility.** Since the datasets only contain the image URLs, and not the images themselves, they are plagued with *link rot* [27].<sup>1</sup> When accessing LAION-5B, there is no guarantee the images still exist at their URLs, making it impossible to fully reproduce the dataset and opening up the possibility of data poisoning attacks [8].

A natural alternative is to not use LAION datasets for training. One could instead independently curate a dataset of CC-licensed images with known provenance that expressly allow for copying, adaptation, and commercial use. As constituent images can be stored and distributed, this would also solve the link rot problem, thereby enabling greater reproducibility. We defer our discussion of sourcing CC-licensed images to Section 4, where we detail *CommonCatalog*: our new, open dataset. While CC images are an attractive alternative to LAION-5B, we note that CC images rarely contain the captions necessary to train T2I models. Therefore, we first need a method for captioning CC images, which we describe in the next section.

### 3 Telephoning: A Transfer Learning-based Image-captioning Method

Our solution for handling the lack of captions in CC images is an intuitive type of transfer learning for producing high-quality synthetic labels. We describe this method, and then note that there are various similar methods in prior generative modeling literature. Altogether, these methods indicate that this type of transfer learning to produce synthetic labels (to later serve as inputs to training other generative models) has become an increasingly common pattern. We therefore give this method a name: *telephoning*.

#### 3.1 Describing telephoning

Telephoning (Figure 3) takes inputs from a high-dimensional modality (e.g., images), effectively performs a “lossy compression” to a low-dimensional modality (e.g., short-text captions), and then

<sup>1</sup>This also applies to other scraped datasets, such as DataComp [14] and OBELICS [28].

decompresses back to the high-dimensional modality. Because the intermediate compression step is “lossy”, the ultimate output often does not remotely resemble the original input, just like a game of telephone [38]. We derive the term telephoning from the above intuition, and employ it as useful shorthand to denote instances of transfer learning that solve data-scarcity problems in multimodal generative modeling.

In this work, CC images are the high-dimensional inputs, and we use a pre-trained BLIP-2 model [34] for “lossy compression” to short-text captions (Figure 3a). Together, these CC-image-caption pairs comprise the CommonCatalog dataset, which we use to train our CommonCanvas T2I models (Figure 3b). Even though BLIP-2 was pre-trained on LAION-400M [52], CommonCatalog and CommonCanvas never have direct access to LAION-400M or, importantly, anything that is similar to the images that BLIP-2 was trained on. Instead, we only have access to the mapping in the model, which, given an image input, produces lossy output text that inherently does not literally resemble its image counterpart (Figure 3c).<sup>2</sup>

We defer to experts about fair use (Section 2.2) — namely, regarding models like BLIP-2, and LAION-5B’s images and alt-text captions. Generally, these experts seem to think that many cases will fall under fair use [29, 32, 50], especially when model outputs do not resemble their inputs, which is the case with BLIP-2.

### 3.2 Related work on telephoning

Our work aligns with the trend of using advanced generative models to address data scarcity. This is evident in various modalities, such as producing audio captions from image-text pairs [64] and text from audio [46]. Similar approaches have also been used to generate instruction tuning datasets for both text and images [35, 37]. Concurrent work has used visual question answers models such as LLaVa [37] to enhance existing captions such as such as DALLE-3 [4] and Chen et al. [9]. However, our model is the one of the first work to train on a dataset without any ground truth captions, and one of the first to release our synthetic captioning dataset along with a fully trained diffusion model. Furthermore, the caption upsampling approaches described in these works could be used to further improve the captions of the CommonCatalogue in future work. Captioning models have been used before to create descriptive captions before to guide a diffusion model to create an image visually similar to a specific image. The concurrent work SynthCap [6] generates a synthetic captioning dataset using a diffusion model to generate images from captions, tackling the inverse of our problem statement.

We coin the term telephoning to shorthand processes like these, which include our work and prior work, and which we believe will become more prevalent as generative models progress.

## 4 CommonCatalog: A Dataset of CC Images & Synthetic Captions

In this section, we introduce our open dataset, *CommonCatalog*. First, we describe the collection and curation process for the open-licensed, CC images. This process brings to light two challenges: caption-data incompleteness and image-data scarcity. To address the lack of CC captions, we show concretely how we use telephoning to produce high-quality synthetic captions to accompany our set of curated images. We investigate the topic of data scarcity in the next section, where we also discuss necessary systems-level training optimizations that enable us efficient SD-model iteration.

### 4.1 Sourcing provenanced, licensed images for CommonCatalog

We focus on locating high-resolution Creative-Commons images that have open licenses. We began with the YFCC100M dataset, which consists of 100 million CC-licensed images and multimedia files, as well as Flickr IDs linking to the original data [59]. The images in the dataset associated with the original paper exhibit two issues that make it ill-suited for direct use to train Stable Diffusion: they are low-resolution, and many of them have licenses that do not expressly allow for the distribution of derivative works, which are an area of unsettled copyright law in the context of model training. We therefore re-scraped these images from Flickr, based on the IDs provided in the YFCC100M metadata. Our scraped images are very high resolution (exceeding 4K), which makes them more suitable for T2I training.

<sup>2</sup>We draw on the example of Snoopy from [49]. Figure 3’s Snoopy is CC-licensed [54].



Source	Caption
Alt-Text (LAION-2B)	Latest 1PC Transparent Gradient Color Voile Window Curtain
BLIP2-OPT-2.7B	A living room with a white couch and curtains

Figure 5: Original vs. BLIP-2-generated captions for an image from LAION-2B. BLIP-2 generates a caption that better aligns with what a human would write. See Figure 14 for more examples.

We exclude images with non-derivative (ND) licenses. The remaining images can be further divided into those that can be used for commercial (C) purposes and those that cannot (non-commercial/ NC). As shown in Table 4, we accordingly construct two datasets, CommonCatalog-C and CommonCatalog-NC. We defer additional details about licenses to Appendix B.1.1, but emphasize that all of the images included have open licenses: individuals are free to use, adapt, and remix the images, so long as they attribute them. In total, CommonCatalog contains roughly 70 million NC CC-images, of which a subset of approximately 25 million images can also be used commercially.

Figure 4: CommonCatalog-C contains images licensed only for commercial use; -NC contains -C as well as images licensed for non-commercial use.

Dataset	# Images	% Alt Text
CommonCatalog-C	26,232,417	30.76%
CommonCatalog-NC	67,015,331	31.22%

Directly sourcing CommonCatalog avoids some concerns (Section 2.2); however, it also comes with its own challenges. For one, CC images rarely have the alt-text captions necessary to train a T2I model like Stable Diffusion (Figure 4); those that do have associated text often just include the image title or a URL. For another, we could *only* find roughly 70 million usable CC images, which pales in comparison to the billions of images in LAION used to train SD2 (Section 5). We take each of these challenges in turn. First, in the next subsection, we show how we instantiate telephoning (Section 3) to produce high-quality, synthetic captions for CC images.

## 4.2 Synthesizing captions with telephoning

We compared several captioning models and, based on qualitative analysis and its state-of-the-art performance on MS COCO, chose to use the pre-trained BLIP-2 OPT2.5B model for synthesizing CommonCatalog’s captions [34]. BLIP-2 consists of three components: a pre-trained, frozen (i.e., fixed) visual encoder, a learned transformer network that converts the visual embeddings into a text prompt, and a frozen large language model (LLM) that takes in the prompt. The only trainable variables in the transformers are between the frozen visual encoder and frozen LLM layers.

Given a LAION-2B image as input, we found that the resulting BLIP-2 caption is often qualitatively more descriptive than the corresponding LAION-2B ground-truth alt-text caption. LAION-2B captions often contain product names, irrelevant details, or poor grammar and syntax (Figure 5). This finding is corroborated by Nguyen et al. [42], which shows quantitatively (in terms of CLIP Score) that BLIP-2 captions are higher quality than ground-truth captions, at the cost of caption diversity.

Based on these preliminary results, we captioned all of the YFCC100M Creative-Commons images, which required about 1,120 GPU A100 hours. To do so, we center-cropped and resized all of the images to a maximum size of 512x512 pixels. We perform these transformations because captioning images at native resolution would be very expensive. At training time of the diffusion model, all images remain in their native resolution. We release our commercial (CommonCatalog-C) and non-commercial (CommonCatalog-NC) CC-image and synthetic-caption datasets on HuggingFace at [REDACTED] with associated data cards. As an evaluation set, we also release the BLIP-2 captions that we produced for the non-derivative (ND) CC images that we did not use for training.

## 5 Training Efficiency Optimizations and Data Scarcity Analysis

High-resolution CC images are indeed much less abundant than arbitrary web-scraped ones, but the amount of data necessary to train high-quality SD2 models has not been well-studied. We set out to quantify this amount by training multiple SD2 models on differently-sized subsets of LAION-2B. However, training a single SD2 model, even with hundreds of GPUs, can take several days. To make our data scarcity analysis more tractable, we first implement several efficiency optimizations.

## 5.1 Software and hardware speed-ups

Stability AI reports an estimated 200,000 A100 hours to train SD2 [58]. Depending on the available hardware, a single SD2 run could take anywhere from a few weeks to over a month to train. We sought out multiple avenues to reduce this training-time constraint. Ultimately we were able to achieve a speedup of 2.71X relative to the original SD2 implementation.

First, we applied Flash Attention [11] with the xFormers library [31]. We also pre-computed VAE and text encoder latents over the entire training dataset, cast all GroupNorm [63] and LayerNorm [2] to float16 precision, and applied fully-sharded data parallelism (FSDP) to our training run. Finally we opted to only keep an exponential moving average of the weights for the final 3.5% of training. More detail on each of these improvements can be found in Appendix D.

When applying all of the aforementioned strategies together, we are able to achieve a 2.71X speedup in A100 hours over our SD2-baseline implementation. We found that latent pre-computation helped the most at low resolutions, while FSDP also provided significant gains, especially at scale. The other optimizations helped reduce total memory usage, allowing us to increase the microbatch size for better hardware utilization. Figure 6 summarizes each of the proposed methods and the cumulative speedup that results from its application. Equipped with an optimized training setup, we are able to more easily study effect of varying training dataset size.

## 5.2 Investigating data scarcity: Saturating SD2 evaluations with < 3% of LAION-2B

YFCC100M contains 100 million images, about 10% the size of the 1.1B LAION examples we could access, thus about 5% of the original LAION-2B dataset. One interesting question that remains unanswered is how much data is actually needed to train these diffusion models effectively.

We ask whether or not it is necessary to train on 1+ billion images to get results that are as good as the original LAION-trained SD2. Our results show, surprisingly, that this is not the case with a slightly larger model (CommonCanvas-L); this model replaces SD2’s U-Net with SDXL’s [43] larger one. Further, our larger model achieves comparable results to SD2-base on human evaluation, using 33X less training data. We train on increasingly smaller, random subsets of data from our LAION-1.1B model and find that we can achieve a similar result on the commonly reported MS COCO numbers, but with <3% the amount of SD2’s training data (Figure 8). In fact, we run experiments down to 1-million LAION-1.1B images, and find that only 10 million images are required for stable training behavior (Appendix, Figure 15).

## 5.3 Investigating the performance of CC trained model

These findings suggest that SD2 models may be underparameterized. In fact, when we use CommonCanvas-LNC, we achieve competitive performance with SD2 on user preferences, despite training on significantly less data (Section 7). Further, in spite of the drastic reduction in dataset size, we observe that the larger model (CommonCanvas-LNC) outperforms the smaller one (CommonCanvas-SNC), consistent with the notion that these models are still underparameterized. We hypothesize about why this might be the case and how much data is actually necessary to saturate the model in Appendix A.1.




Figure 6: Cumulative effect of various speed-ups in our SD2 training pipeline on 128 Throughputs evaluated on 128 A100s.




Figure 7: User preference study using Parti prompts. CommonCanvas-LNC model matches the performance of SD2 despite being trained with < 3% the amount of data.




Figure 8: FID, KID, and CLIP-FID vs. CLIP-Score computed on 30K samples from COCO2014 for different SD2 models trained on smaller subsets of LAION (10M, 90M, using either original captions or synthetic BLIP2 captions). Interestingly, increasing the amount of training data from 10M to 90M samples does not lead to improved quantitative metrics across guidance scales 1 to 8. Lower FID is better; higher CLIP score is better.

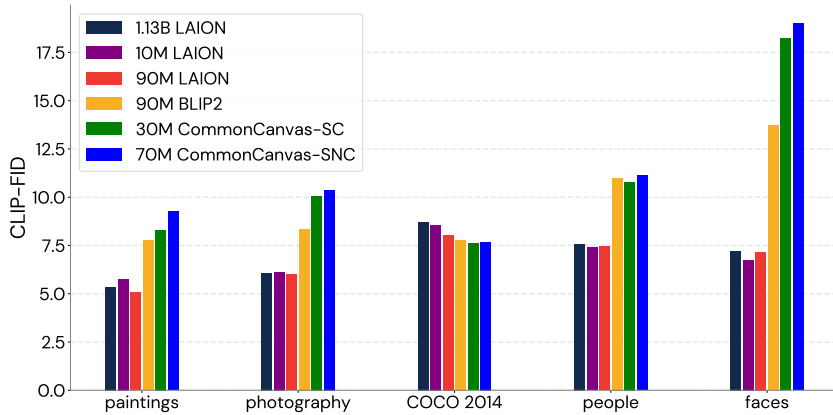


Figure 9: CLIP-FID for different models. We can see domain shift between MS COCO captions and web-scraped conceptual captions. CLIP-FID likely favors SD2, as CLIP is trained on a similar style of text as LAION. This plot only covers the first stage of training at 256x256 resolution. We

## 6 Experiments

Equipped with commercial (CommonCatalog-C) and non-commercial (CommonCatalog-NC) datasets, we train two different CommonCanvas models. We additionally train a larger variant of CommonCanvas-NC (CommonCanvas-LNC) that, as we note above (Section 5.2), has a significantly larger U-Net. Figure 1 displays qualitative results from each of these model variants. More details on the CommonCanvas-LNC architecture can be found in Appendix A.2.

### 6.1 Automated quality metrics for model evaluation

We measure performance with three automated image quality metrics on the MS COCO dataset [36]: Frechet Inception Distance (FID) [17], Kernal Inception Distance (KID) [5], and CLIP-FID [25]. Additionally, CLIP Score was evaluated to understand the alignment between captions and their respective images. Our model demonstrated comparable performance compared to the baseline of SD2 on the popular MS COCO benchmark.

However, like any model, ours has limitations. It underperformed in several categories, including faces, general photography, and paintings. These categories originated from the Conceptual Captions dataset [55], which relies on web-scraped data. These web-sourced captions, while abundant, may not always align with human-generated language nuances.

This discrepancy underscores the importance of incorporating large-scale, human-generated caption data. Although transitioning to synthetic captions introduces certain performance challenges, the





Figure 10: We compare CommonCanvas-SNC (Ours) to SD2. Our model is less likely to generate iconic characters given suggestive prompts (drawn from Lee et al. [29]).

drop in performance is not as dramatic as one might assume. Moreover, we speculate that it would if users were to supplement with their own datasets, like FFHQ [22], if they seek to fine-tune models for specific categories.

## 6.2 Human evaluation

While automated quality metrics are useful, given the level of detail and breadth of the distribution large T2I are intended to generate, there is no substitute for evaluation by human raters. Human pairwise preference ratings for the three 512x512 resolution CommonCanvas models compared to SD2-base can be seen in Figure 7.

In this experiment, human raters were shown a prompt (selected randomly from the PartiPrompts prompts set [66]) along with two generated images in randomized order, one from the reference model (SD2-base) and the other from a CommonCanvas model. Users were asked which generated image they preferred. We report the fraction of the time users selected the image generated by the CommonCanvas model over the corresponding generation from SD2 as the user preference rate for that model. In agreement with our automated quality metrics, we find that the two small CommonCanvas models are less preferred than SD2-base, with preference rates of 37% for CommonCanvas-SC and 38% for CommonCanvas-SNC, which we find surprisingly high considering the smaller and synthetic nature of the dataset. For the largest model, CommonCanvas-LNC, we do not measure a statistically significant difference in user preference between this model and SD2-base. While SDXL is a significantly larger model, this finding represents an existential result, showing that we are capable of matching the performance of a model trained on several magnitudes more of data.

## 6.3 Benefits and challenges of synthetic captions

Interestingly, we observe that synthetic captions can enhance the alignment of our model. For instance, the CLIP Score for synthetic captions exceeded that of ground-truth captions as seen in Figure 8.

We also observed reduced diversity of n-grams in our synthetic captions, a pattern previously noted by Nguyen et al. [42]. This effect can be visualized through the decrease in unique trigrams.

Although we train on Creative-Commons images, it is still possible for an adversarial prompt to produce content that, for example, includes iconic characters. In Figure 10, we subject our model to ambiguous prompts that are suggestive of such characters. Examples include visuals closely resembling Elsa from Frozen, Indiana Jones resembling Harrison Ford, and even a likeness to Harry Potter (Figure 10). Qualitatively, our model deviated more from these characters than SD2.



Figure 11: Using CommonCanvas-SNC (Ours) to generate celebrities. Our model is worse at synthesizing individual people than SD2, but is capable of generating some noteworthy public figures.

## 7 Discussion and Related Work

In this paper, we train the family of CommonCanvas text-to-image latent diffusion models on only Creative-Commons images and synthetic captions. We discuss the data incompleteness and scarcity issues associated with CC images, and how we address each of these issues in turn. For data incompleteness, we propose telephoning, an intuitive type of transfer learning (Section 3), which we instantiate with BLIP-2 to produce synthetic captions for CC images — together, the CommonCatalog dataset (Section 4). With regard to data scarcity, we hypothesize that much less data than what is contained in LAION-2B is necessary to saturate SD2, and that CommonCatalog should be sufficient for training. To make testing this hypothesis more efficient, we implement a variety of ML-systems optimizations, which achieve a 2.7X speed-up over our SD2 baseline. Ultimately, we find that we can train SD2 on <3% of LAION-2B (Section 5), which encourages us to train on CommonCatalog’s commercial (roughly 70 million) and non-commercial (roughly 25 million) examples. Our CommonCanvas models under-perform in some categories, like faces, but CommonCanvas-LNC demonstrates statistically equivalent performance with SD2 on human evaluation (Section 6).

We note that several recent works study copyright. This work tends to concern text-to-text training data [39], be primarily theoretical [51, 62], involve ablation studies [24], or only handle verbatim memorization [7] through the use of generation-time content filters [16], which has been shown to be an incomplete solution [19]. To the best of our knowledge, no prior open work attempts to train T2I models on only open-licensed data.

Most prior work on text-caption-dataset creation has focused on extracting caption data from Common Crawl [12, 14, 28]. We instead focus on synthesizing captions directly by using a pre-trained BLIP-2 model. [42] demonstrate that existing caption datasets can be improved by using BLIP2 to re-caption low-quality captions in large datasets like Datacomp, but do not focus on creating a new dataset of synthetic captions, as we do here.

An issue, which we do not address, is that the YFCC100M data is about a decade old; its CC images are not as current as those in LAION-2B. Given the success of our results, in the future, we plan to augment CommonCatalog with Creative-Commons images from other sources, as well as test larger CommonCanvas model architectures.

## Acknowledgements

We would like to thank Christopher De Sa for feedback on earlier drafts of this work. A. Feder Cooper is funded by Professor Christopher De Sa’s NSF RI-CAREER award 2046760. This work was also sponsored by Volodymyr Kuleshov’s CAREER grant: #2145577. We also would like to thank Apolinário Passos for helping us host the data + models and for insightful discussions along the way.

## References

- [1] Anderson v. Stability AI, Ltd., 2023. No. 3:23-cv-00201 (N.D. Cal. Jan. 13, 2023).
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- [3] Romain Beaumont. LAION-5B: A New Era of Large-Scale Multi-Modal Datasets. *LAION Blog*, March 2022. URL <https://laion.ai/blog/laion-5b/>.
- [4] James Betker, Gabriel Goh, Li Jing, Tim Brooks, Jianfeng Wang, Linjie Li, Long Ouyang, Juntang Zhuang, Joyce Lee, Yufei Guo, Wesam Manassra, Prafulla Dhariwal, Casey Chu, Yunxin Jiao, and Aditya Ramesh. Improving image generation with better captions. 2023.
- [5] Mikołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. *arXiv preprint arXiv:1801.01401*, 2018.
- [6] Davide Caffagni, Manuele Barraco, Marcella Cornia, Lorenzo Baraldi, and Rita Cucchiara. Synthcap: Augmenting transformers with synthetic data for image captioning. In *International Conference on Image Analysis and Processing*, pages 112–123. Springer, 2023.
- [7] Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. Extracting Training Data from Large Language Models. In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association, August 2021. ISBN 978-1-939133-24-3.
- [8] Nicholas Carlini, Matthew Jagielski, Christopher A. Choquette-Choo, Daniel Paleka, Will Pearce, Hyrum Anderson, Andreas Terzis, Kurt Thomas, and Florian Tramèr. Poisoning Web-Scale Training Datasets is Practical, 2023.
- [9] Junsong Chen, Jincheng Yu, Chongjian Ge, Lewei Yao, Enze Xie, Yue Wu, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- $\alpha$ : Fast training of diffusion transformer for photorealistic text-to-image synthesis, 2023.
- [10] Mehdi Cherti, Romain Beaumont, Ross Wightman, Mitchell Wortsman, Gabriel Ilharco, Cade Gordon, Christoph Schuhmann, Ludwig Schmidt, and Jenia Jitsev. Reproducible scaling laws for contrastive language-image learning, 2022.
- [11] Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. FlashAttention: Fast and memory-efficient exact attention with IO-awareness. In *Advances in Neural Information Processing Systems*, 2022.
- [12] Karan Desai, Gaurav Kaul, Zubin Aysola, and Justin Johnson. Redcaps: Web-curated image-text data created by the people, for the people. *arXiv preprint arXiv:2111.11431*, 2021.
- [13] Doe 1 v. GitHub, Inc., 2022. No. 4:22-cv-06823 (N.D. Cal. November 3, 2022).
- [14] Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, et al. DataComp: In search of the next generation of multimodal datasets, 2023.
- [15] Getty Images (US), Inc. v. Stability AI, Inc., 2023. No. 1:23-cv-00135 (D. Del. February 3, 2023).
- [16] GitHub. Configuring github copilot in your environment, 2023. URL <https://docs.github.com/en/copilot/configuring-github-copilot/configuring-github-copilot-in-your-environment>.
- [17] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *arXiv preprint arXiv:2104.08718*, 2021.
- [18] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. OpenCLIP, July 2021. URL <https://doi.org/10.5281/zenodo.5143773>.
- [19] Daphne Ippolito, Florian Tramèr, Milad Nasr, Chiyuan Zhang, Matthew Jagielski, Katherine Lee, Christopher A. Choquette-Choo, and Nicholas Carlini. Preventing Verbatim Memorization in Language Models Gives a False Sense of Privacy, 2023.
- [20] J.L. v. Alphabet Inc., 2023. No. 3:23-cv-03440-LB (N.D. Cal July 11, 2023).

- [21] Kadrey v. Meta Platforms, Inc., 2023. No. 3:23-cv-03417 (N.D. Cal. July 7, 2023).
- [22] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4401–4410, 2019.
- [23] Dirk P. Kingma and Max Welling. Auto-Encoding Variational Bayes. In *International Conference on Learning Representations*, 2014.
- [24] Nupur Kumari, Bingliang Zhang, Sheng-Yu Wang, Eli Shechtman, Richard Zhang, and Jun-Yan Zhu. Ablating Concepts in Text-to-Image Diffusion Models, 2023.
- [25] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fr`echet inception distance. *arXiv preprint arXiv:2203.06026*, 2022.
- [26] LAION-2Ben, 2022. URL <https://huggingface.co/datasets/laion/laion2B-en>. Accessed September 23, 2023.
- [27] Viktor Lakic, Luca Rossetto, and Abraham Bernstein. Link-Rot In Web-Sourced Multimedia Datasets. In *MultiMedia Modeling: 29th International Conference, MMM 2023, Bergen, Norway, January 9–12, 2023, Proceedings, Part I*, page 476–488, Berlin, Heidelberg, 2023. Springer-Verlag. ISBN 978-3-031-27076-5. doi: 10.1007/978-3-031-27077-2\_37. URL [https://doi.org/10.1007/978-3-031-27077-2\\_37](https://doi.org/10.1007/978-3-031-27077-2_37).
- [28] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. OBELICS: An Open Web-Scale Filtered Dataset of Interleaved Image-Text Documents, 2023.
- [29] Katherine Lee, A. Feder Cooper, and James Grimmelmann. Talkin’ `Bout AI Generation: Copyright and the Generative-AI Supply Chain, 2023.
- [30] Katherine Lee, A. Feder Cooper, James Grimmelmann, and Daphne Ippolito. AI and Law: The Next Generation, 2023.
- [31] Benjamin Lefaudeux, Francisco Massa, Diana Liskovich, Wenhan Xiong, Vittorio Caggiano, Sean Naren, Min Xu, Jieru Hu, Marta Tintore, Susan Zhang, Patrick Labatut, and Daniel Haziza. xFormers: A modular and hackable Transformer modelling library. <https://github.com/facebookresearch/xformers>, 2022.
- [32] Mark A. Lemley. How Generative AI Turns Copyright Law on its Head, 2023. URL <https://ssrn.com/abstract=4517702> or <http://dx.doi.org/10.2139/ssrn.4517702>.
- [33] Pierre N. Leval. Toward a Fair Use Standard. *Harvard Law Review*, 103(5):1105, 1990.
- [34] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [35] Xian Li, Ping Yu, Chunting Zhou, Timo Schick, Luke Zettlemoyer, Omer Levy, Jason Weston, and Mike Lewis. Self-alignment with instruction backtranslation. *arXiv preprint arXiv:2308.06259*, 2023.
- [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [37] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [38] Susan Box Mann. The Telephone Game, 2019. URL <https://icebreakerideas.com/telephone-game/>. Accessed September 27, 2023.
- [39] Sewon Min, Suchin Gururangan, Eric Wallace, Hannaneh Hajishirzi, Noah A. Smith, and Luke Zettlemoyer. SILO Language Models: Isolating Legal Risk In a Nonparametric Datastore, 2023.
- [40] The Mosaic ML Team. composer. <https://github.com/mosaicml/composer/>, 2021.
- [41] The Mosaic ML Team. streaming. <https://github.com/mosaicml/streaming/>, 2022.

- [42] Thao Nguyen, Samir Yitzhak Gadre, Gabriel Ilharco, Sewoong Oh, and Ludwig Schmidt. Improving multimodal datasets with image captioning. *arXiv preprint arXiv:2307.10350*, 2023.
- [43] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis, 2023.
- [44] Jacob Portes, Alexander R Trott, Sam Havens, Daniel King, Abhinav Venigalla, Moin Nadeem, Nikhil Sardana, Daya Khudia, and Jonathan Frankle. Mosaicbert: How to train bert with a lunch money budget. In *Workshop on Efficient Systems for Foundation Models@ ICML2023*, 2023.
- [45] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- [46] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pages 28492–28518. PMLR, 2023.
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-Resolution Image Synthesis with Latent Diffusion Models. In *2022 IEEE Conference on Computer Vision and Pattern Recognition*, 2022.
- [48] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional Networks for Biomedical Image Segmentation. *Medical Image Computing and Computer-Assisted Intervention*, pages 234–241, 2015.
- [49] Matthew Sag. Copyright Safety for Generative AI. *Houston Law Review*, 2023. Forthcoming.
- [50] Pamela Samuelson. Generative AI meets copyright. *Science*, 381(6654):158–161, 2023. doi: 10.1126/science.adi0656. URL <https://www.science.org/doi/abs/10.1126/science.adi0656>.
- [51] Sarah Scheffler, Eran Tromer, and Mayank Varia. Formalizing Human Ingenuity: A Quantitative Framework for Copyright Law’s Substantial Similarity. In *Proceedings of the Symposium on Computer Science and Law*, pages 37–49, 2022.
- [52] Christoph Schuhmann. LAION-400-Million Open Dataset, 2021. URL <https://laion.ai/blog/laion-400-open-dataset/>.
- [53] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. LAION-5B: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022.
- [54] Charles M. Schultz. Snoopy\_Peanuts, 2020. URL [https://en.wikipedia.org/wiki/Snoopy#/media/File:Snoopy\\_Peanuts.png](https://en.wikipedia.org/wiki/Snoopy#/media/File:Snoopy_Peanuts.png). Accessed September 26, 2023.
- [55] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2556–2565, 2018.
- [56] Benjamin L.W. Sobel. Artificial Intelligence’s Fair Use Crisis. *Columbia Journal of Law and The Arts*, 41:45, 2017.
- [57] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathany, and Surya Ganguli. Deep Unsupervised Learning using Nonequilibrium Thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, 2015.
- [58] Stability AI. Stable Diffusion v2-base Model Card, 2022. URL <https://huggingface.co/stabilityai/stable-diffusion-2-base>.
- [59] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- [60] Tremblay v. OpenAI, Inc., 2023. No. 3:23-cv-03223 (N.D. Cal. June 28, 2023).

- [61] James Vincent. Getty Images is suing the creators of AI art tool Stable Diffusion for scraping its content. *The Verge*, January 2023. URL <https://www.theverge.com/2023/1/17/23558516/ai-art-copyright-stable-diffusion-getty-images-lawsuit>.
- [62] Nikhil Vyas, Sham Kakade, and Boaz Barak. On Provable Copyright Protection for Generative Models, 2023.
- [63] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018.
- [64] Feiyang Xiao, Qiaoxi Zhu, Jian Guan, Xubo Liu, Haohe Liu, Kejia Zhang, and Wenwu Wang. Synth-ac: Enhancing audio captioning with synthetic supervision. *arXiv preprint arXiv:2309.09705*, 2023.
- [65] Yuanzhong Xu, HyoukJoong Lee, Dehao Chen, Hongjun Choi, Blake Hechtman, and Shibo Wang. Automatic cross-replica sharding of weight update in data-parallel training. *arXiv preprint arXiv:2004.13336*, 2020.
- [66] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2(3):5, 2022.

## A Additional Details on Data Scarcity Analysis

### A.1 Hypothesis: Diffusion models are too small

A back-of-the-envelope calculation provides some insight on why this is the case. Consider a training dataset consisting of  $N$  images with resolution  $H \times W$  and  $c$  channels. To completely memorize the training data, the model must be capable of storing  $c \times H \times W \times N$  numbers. Given a number of trainable parameters  $N_p$ , it is natural to assume that on average each parameter is capable of storing roughly enough information to reconstruct a single number from the training dataset. Under this assumption, complete memorization is only possible if the size of the training dataset is at or below a critical size  $N_c$  ( $N \leq N_c$ ) with  $N_c$  given by  $N_c = \frac{N_p}{cHW}$ . Note that this critical size assumes the data cannot be further compressed, which is obviously not the case for natural images. However, SD2 and SDXL are latent diffusion models, which first use a pretrained encoder to compress images by a factor of 8 in both  $H$  and  $W$ , and so when we train LDMS like SD2 and SDXL, we are training on data that has been significantly compressed already.

In our experiments,  $c = 4$  and  $H = W = 32$ , corresponding to  $256 \times 256$  resolution RGB images in the SD2 and SDXL latent space. The SD2 UNet has  $N_p = 866 \times 10^6$  trainable parameters, and SDXL’s UNet has  $N_p = 2567 \times 10^6$ . So we calculate  $N_c \approx 0.2 \times 10^6$  for SD2 and  $N_c \approx 0.6 \times 10^6$  for CommonCanvas-Large; both of these numbers are several orders of magnitude below the size of our YFCC derived datasets, and so even with significant additional data compression we expect that our CommonCatalog datasets should be sufficient to train both SD2 and SDXL. Additionally, this argument predicts that we should only begin to see significant overfitting in these models for datasets of size  $N \sim 10^6$ . These estimates are resolution dependent, and as image resolution increases we expect that  $N_c$  will decrease as more information is provided per image.

### A.2 Increasing model capacity with CommonCanvas-LNC

We also train a variant of SD2 with more trainable parameters, taking the UNet from SDXL. We refer to this model as CommonCanvas-LNC. We adapt the SDXL UNet architecture to SD2 by changing the cross-attention dimensionality to match that of the SD2 text encoder hidden state dimensionality (1024 for SD2 vs. 2048 for SDXL). SDXL also retrains the VAE component in their model, and we use this improved performance VAE as well. Except for these changes, the architecture is identical to that of SD2.

## B Training Dataset Details

### B.1 LAION-2B

The fact that LAION is not a stable benchmark can lead to multiple reproducibility and security issues. Data poisoning attacks would be difficult to detect at the scale of 2 billion parameters. While this could be mitigated by using hash values of the images, then any time the a site decide to re-encode the image, those images would now need to be excluded from the dataset. Furthermore, targeted data poisoning attacks for diffusion models are no longer just academic conjecture. Last year after the release of Stable Diffusion, a protest was launched on ArtStation that had uses upload images that said “NoAI” to taint future training data for generative models after artists felt as though their work had been unfairly used to train the models. With the high degree of link rot, targeted attacks are fairly easy. Furthermore, reproduction of the experiments becomes virtually impossible. This means any benchmarks that use copies of LAION as ground truth are are likely using differing subsets of the full dataset.

#### B.1.1 Sourcing Creative-Commons images

Table 1: CC licenses in YFCC100M. ND means derivative works are not licensed or the license doesn’t allow the user to create derivative works. NC means images cannot be used in commercial contexts. CommonCatalog-C only contains data from the bottom two (yellow) rows, reflecting images licensed for commercial contexts (i.e., roughly 25 million images). CommonCatalog-NC contains CommonCatalog-C, and additionally includes the middle two (blue) rows, reflecting images licensed for non-commercial purposes. We do not include the roughly 30 million images in the top two (pink) rows in CommonCatalog, as they are non-derivative licenses. We do not train on these images. We do, however, produce BLIP-2 captions for them and release those captions as an evaluation set.

CC License	# Images	% Captioned
CC-BY-NC-ND-2.0	25,790,117	33.52%
CC-BY-ND-2.0	4,827,970	30.23%
CC-BY-NC-2.0	12,468,229	31.39%
CC-BY-NC-SA-2.0	28,314,685	31.57%
CC-BY-SA 2.0	9,270,079	34.05%
CC-BY 2.0	16,962,338	28.96%


#### B.1.2 Release and documentation

## C YFCC Example Images

Table 2: Randomly sampled images from the YFCC [59] training set. Our synthetic BLIP2 captions are also provided below.



a person riding a bike on a dirt road



a paintings on the wall



an orange and blue race car driving on a track

## Model Architecture

Table 3: Top 10 highest frequency captions in the YFCC dataset. The most common captions are not user generated and are not very descriptive of the corresponding image.

YFCC Original Caption	Count
OLYMPUS+DIGITAL+CAMERA	184889
SONY+DSC	123128
Exif_JPEG_PICTURE	104480
Barclays+Center+Arena%0AAtlantic+Yards%0A6th+and+Atlantic+A	68832
Olympus+digital+camera	54805
Effortlessly+uploaded+by Eye-Fi	48388
.	43227
-+Camera+phone+upload+powered+by ShoZu	38856
Sony+dsc	32709
Photo+by @Kmeron —Facebook page is this way—	23754

Table 4: Number of usable captions from OpenAI’s YFCC14M dataset [45]. This table is actually a subset from 1 for which either the user description or image title were deemed usable. These figures provide an estimate on how many images in each category are actually potentially usable as captions.

License Name	count
CC-BY 2.0	2448002
CC-BY-ND 2.0	682273
CC-BY-NC 2.0	1925854
CC-BY-NC-ND 2.0	4058817
CC-BY-NC-SA 2.0	4146113
CC-BY-SA 2.0	1568336

We follow the model architecture and training recipe of Stable Diffusion 2 as closely as we can to best reproduce the model for CC-Small. The model has an identical number of params and structure as the original model. In fact, we can even load SD2’s model weights into our framework due to the identical architecture and naming scheme. We are able to achieve virtually identical performance with SD2 in a much shorter training time with less data. We use the same VAE, tokenizers, and UNet architecture as SD2 except for reducing the precision of the normalization layers.

Our CC-Large model takes SD2’s model and replaces the UNet with the SDXL architecture [43]. Like CC-Small, we also replace the normalization layers with their low-precision version. The replacement of all the normalization layers is handled automatically by MosaicML’s Composer library [40]. We perform all data loading through MosaicML’s streaming library [41].

## D Additional Details on Efficiency Optimizations

In this section we provide additional details on the optimizations we implemented to achieve SD2 training speedups. We also report the approximate cost of training our implementation of SD2 on various hardware configurations in Table 5.

**Flash Attention.** Cross attention operations are a very expensive part of training that occurs in dozens of layers in diffusion model UNets [47]. Flash Attention is an efficient implementation that

Table 5: Performance (throughput) and approximate cost of training SD2 UNet with our optimizations. Depending on the number of GPUs used, the cost to train the same models without these optimizations range from \$90,000-\$140,000

Number of A100s	256x256 (img/s)	512x512 (img/s)	512x512 with EMA (img/s)	Days to Train	Cost (\$)
8	1100	290	290	101.04	\$38,800.00
16	2180	585	580	50.29	\$38,630.00
32	4080	1195	1160	25.01	\$38,420.00
64	8530	2340	2220	12.63	\$38,800.00
128	11600	4590	3927	6.79	\$41,710.00



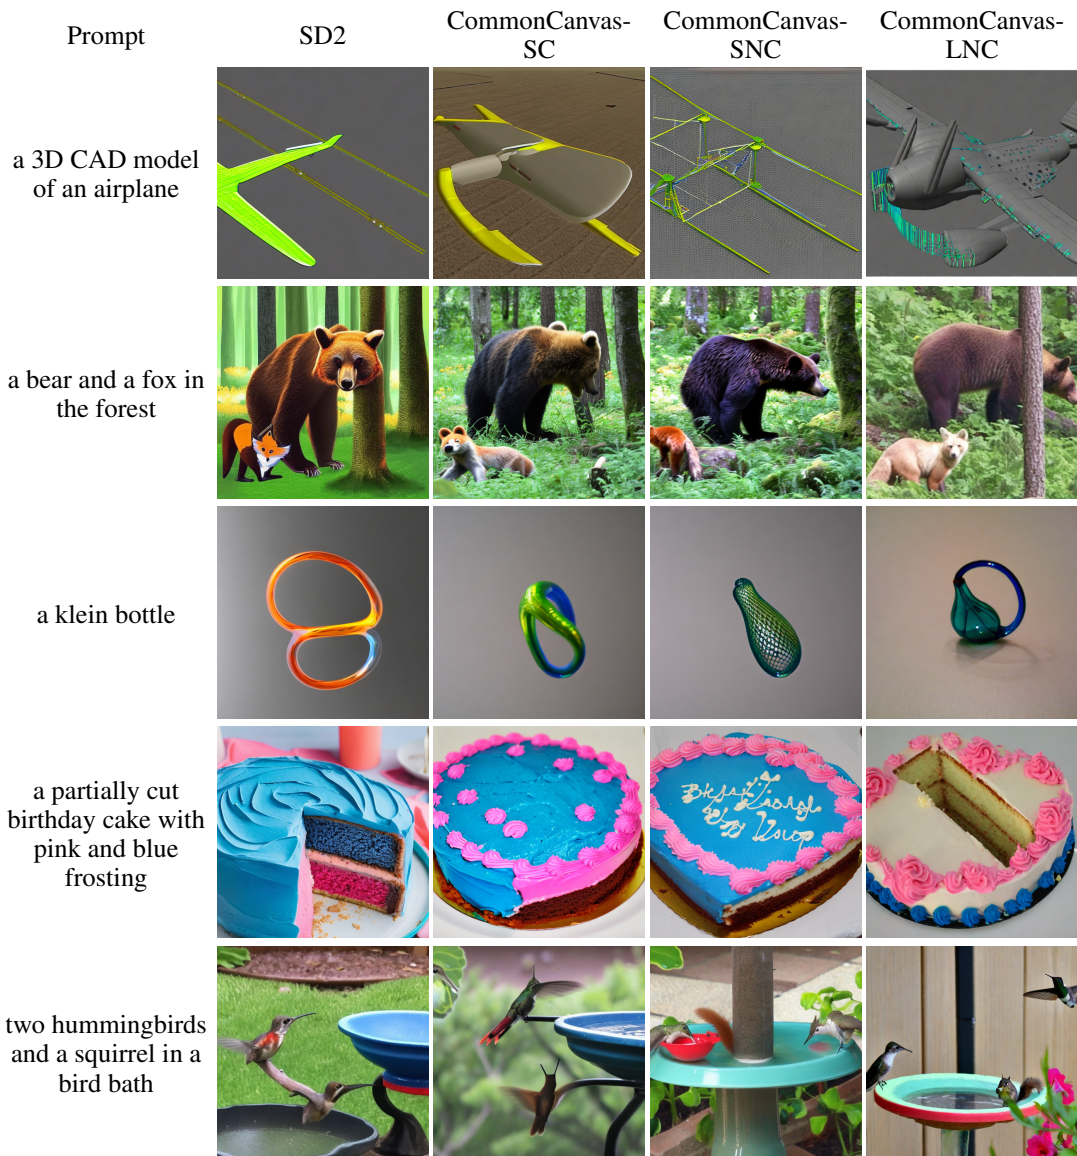


Figure 12: Additional qualitative examples comparing SD2 to our model trained on the commercial split (CommonCanvas-SC), non-commercial split (CommonCanvas-SNC), and the larger UNet model trained on the non-commercial (CommonCanvas-LNC).




Figure 13: Additional qualitative examples of our CommonCanvas models.



Figure 14: Additional qualitative examples comparing our CommonCanvas models to SD2, given synthetic BLIP2 captions as prompts. While not perfect, our models are better at avoiding generating potentially problematic data.

is optimized to work well with reduced precision and GPU hardware [11], which was implemented using the XFormers library [31], allowing us to save compute and memory usage.

**Precomputing latents.** Each forward pass of SD2 requires computing a latent representation of the input image, as well as transforming the caption into a text embedding. Instead of computing the latents for each example during training, we can precompute latents for the entire dataset, amortizing the cost. Doing so speeds up training of the model, especially at lower resolutions, in exchange for a one-time fixed cost of precomputing all the latents over 1 epoch.

**Reduced-precision GroupNorm and LayerNorm.** Most layers in SD2 are implemented in float16 precision, but GroupNorm and LayerNorm are implemented in float32, in part because it was assumed to be necessary for training stability. The resulting, frequent upcasting causes a major bottleneck in training speed. Recent work shows that it is safe to implement LayerNorm using float16 precision [44], and we found the same to be true of GroupNorm. We thus cast all GroupNorm and LayerNorm operators to float16 and are able to further reduce total memory consumption and accelerate training.

**Fully-Sharded Data Parallelism (FSDP).** FSDP is a variant of data-parallel training that shards the models parameters, gradients and optimizer state across multiple devices. When training data batches do not fit into memory, we do several forward and backward passes on smaller microbatches, followed by a single gradient update. At GPU scale, there may only be a single microbatch, so the time for the gradient update can become a significant bottleneck. In standard data distributed training, each GPU communicates all its gradients to every other GPU, and then each GPU updates its local copy of the model. Instead, we use a different paradigm inspired by [65] where each GPU only gets the gradients and updates the weights for a small part of the model before sending the updated weights for that part of the model to all of the other GPUs. By dividing the update step across all the GPUs, we can ensure that the amount of work per GPU decreases as we increase the number of GPUs, helping us achieve linear scaling. To tackle this problem, we use PyTorch’s experimental support for Fully Sharded Data Parallelism (FSDP), specifically, FSDP’s SHARD\_GRAD\_OP mode.




Figure 15: How does reducing the amount of training data affect the training dynamics? We find a noticeable improvement drop when training with less than 10 million samples.

**Scheduled Exponential Moving Average (EMA).** SD2 uses EMA, which maintains an exponential moving average of the weights at every gradient update for the entire training period. This can be slow due to the memory operations required to read and write all the weights at every step. Since the old weights are decayed by a factor of 0.9999 at every batch, the early iterations of training only contribute minimally to the final average. We decide to only apply EMA for the final 50K steps (about 3.5% of the training period), and are able to avoid adding overhead and still achieve a nearly equivalent EMA model.

## E Telephoning: A Transfer Learning-based Image-captioning Method

Our solution for handling the lack of captions in CC images is called *telephoning*, a type of transfer learning (Figure 3). Telephoning assumes the existence of a large labeled dataset  $\mathcal{D}_1 = \{(x^{(i)}, y^{(i)})\}_{i=1}^n$ , consisting of pairs of high-dimensional  $x^{(i)}$  (e.g., images, audio) that map to a compact, structured label  $y^{(i)}$  (e.g., caption, audio transcript). Telephoning trains a forward model  $q(y|x)$  on  $\mathcal{D}_1$  to learn the mapping of  $y$  given  $x$  via maximum likelihood learning  $\max_{q \in \mathcal{Q}} \sum_{i=1}^n \log q(y^{(i)}|x^{(i)})$ . It then uses  $q$  as training signal for a reverse model  $p(x|y)$  trained on a separate dataset  $\mathcal{D}_2 = \{x^{(i)}\}_{i=1}^m$  by maximizing  $\sum_{i=1}^m \mathbb{E}_{y \sim q(y|x^{(i)})} [\log p(x^{(i)}|y^{(i)})]$ , the likelihood of the data  $\mathcal{D}_2$  and the predicted label  $y$  under  $q$ . This forms a type of knowledge transfer from the forward labeling task defined by  $\mathcal{D}_1$  to the reverse task of inverting  $x$  from  $y$  on a separate  $\mathcal{D}_2$ .

While telephoning can be viewed as a type of synthetic labeling, it becomes particularly interesting when  $x$  is a type of protected modality (e.g., a copyrighted image), while  $y$  is a compact representation of  $x$  that does not encode sensitive aspects of  $y$  (e.g., a generic caption). Effectively, telephoning performs a type of “lossy compression” or “distillation” from a high-dimensional or information-rich  $x$  (e.g., an image of Snoopy) to a low-dimensional or information-poor  $y$  that loses the sensitive content in  $x$  (e.g., the visual characteristics of Snoopy). Because this compression step is “lossy”, a reconstruction  $x'$  of  $x$  from  $p(x|y)$  via  $y$  often does not remotely resemble the original input, just like in a game of telephone [38]. We derive the term telephoning from the above intuition, and employ it as useful shorthand to denote instances of transfer learning that solve data-scarcity problems in multimodal generative modeling.

**Telephoning for text-to-image modeling.** In this work, we apply telephoning to the image and text domains, where CC images are the high-dimensional inputs  $x$ , and we use a pre-trained BLIP-2 model [34] for “lossy compression” to short-text captions  $y$  (Figure 3a). Together, these CC-image-caption pairs comprise the CommonCatalog dataset, which we use to train our CommonCanvas T2I models (Figure 3b). Even though BLIP-2 was pre-trained on LAION-400M [52], CommonCatalog and CommonCanvas never have direct access to LAION-400M or, importantly, anything that is similar to the images that BLIP-2 was trained on. Instead, we only have access to the mapping in the model, which, given an image input, produces lossy output text that inherently does not literally resemble its image counterpart (Figure 3c).<sup>3</sup>

<sup>3</sup>We draw on the example of Snoopy from [49]. Figure 3’s Snoopy is CC-licensed [54].