

Visual Atoms: Pre-training Vision Transformers with Sinusoidal Waves

Sora Takashima^{1,2}, Ryo Hayamizu¹, Nakamasa Inoue^{1,2}, Hirokatsu Kataoka¹, Rio Yokota^{1,2}

¹National Institute of Advanced Industrial Science and Technology (AIST)

²Tokyo Institute of Technology

<https://masoral030.github.io/Visual-Atoms-Pre-training-Vision-Transformers-with-Sinusoidal-Waves/>

Abstract

Formula-driven supervised learning (FDSL) has been shown to be an effective method for pre-training vision transformers, where ExFractalDB-21k was shown to exceed the pre-training effect of ImageNet-21k. These studies also indicate that contours mattered more than textures when pre-training vision transformers. However, the lack of a systematic investigation as to why these contour-oriented synthetic datasets can achieve the same accuracy as real datasets leaves much room for skepticism. In the present work, we develop a novel methodology based on circular harmonics for systematically investigating the design space of contour-oriented synthetic datasets. This allows us to efficiently search the optimal range of FDSL parameters and maximize the variety of synthetic images in the dataset, which we found to be a critical factor. When the resulting new dataset VisualAtom-21k is used for pre-training ViT-Base, the top-1 accuracy reached 83.7% when fine-tuning on ImageNet-1k. This is only 0.5% difference from the top-1 accuracy (84.2%) achieved by the JFT-300M pre-training, even though the scale of images is 1/14. Unlike JFT-300M which is a static dataset, the quality of synthetic datasets will continue to improve, and the current work is a testament to this possibility. FDSL is also free of the common issues associated with real images, e.g. privacy/copyright issues, labeling costs/errors, and ethical biases.

1. Introduction

Vision transformers [10] have made a significant impact on the entire field of computer vision, and state of the art models in classification [37, 41, 42], object detection [25, 36], and segmentation [8, 15, 23, 36] are now based on vision transformers. The accuracy of vision transformers exceeds that of convolutional neural networks by a considerable margin when the model is pre-trained on huge datasets, such as JFT-300M [32]. However, the JFT-300M dataset contains 300M images and 375M labels. It is impos-

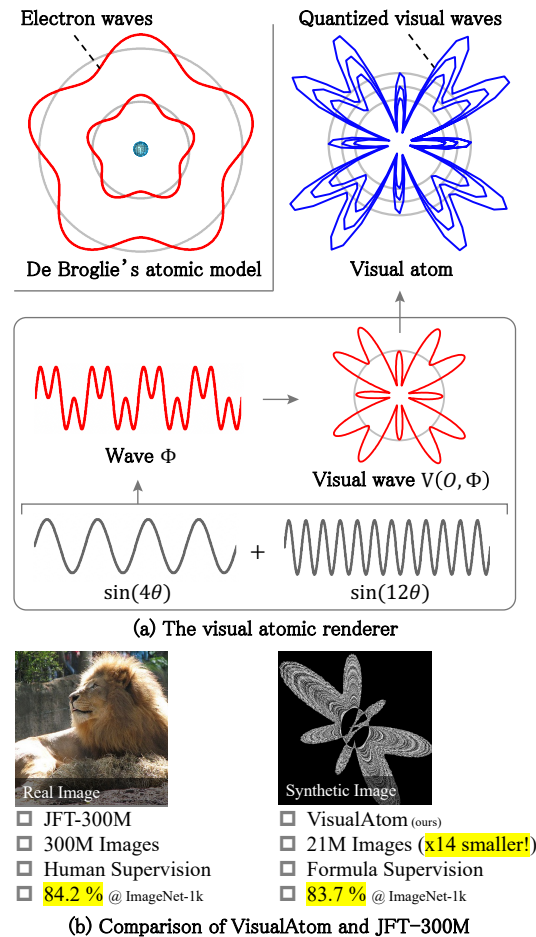


Figure 1. **VisualAtom: a new FDSL dataset** (a) Inspired by de Broglie’s atomic model, we propose VisualAtom dataset containing shapes from two sinusoidal waves. (b) When fine-tuning on ImageNet-1k, ViT-B pre-trained with VisualAtom achieved the accuracy of only 0.5% lower than JFT-300M using 1/14 images.

sible to manually label all of these images. Efforts to automatically label such datasets is still not as accurate as manual labeling. Self-supervised learning (SSL) is increasing in popularity, as datasets do not need to be labeled for this

mode of training [16]. Although SSL removes the burden of labeling large datasets, the effort to collect/download, store, and load these large datasets remains a challenge.

One of the major issues in computer vision is that access to huge datasets, such as JFT-300M/3B [32, 42] and IG-3.5B [27], is limited to certain institutions. This makes it difficult for the rest of the community to build upon, or even reproduce, existing works. This limitation has prompted the creation of open datasets, such as LAION-5B [30]. However, the LAION dataset is not curated, which means that it could contain inappropriate content, or be subject to societal bias and/or privacy/copyright issues [5, 39, 40]. How to curate such large datasets from the perspective of AI ethics and safety is an open area of research, but, in the meantime, an alternative approach to creating large datasets for computer vision is needed.

Formula-Driven Supervised Learning (FDSL) [19, 20] has been proposed as an alternative to supervised learning (SL) and SSL with real images. The term “formula-driven” encompasses a variety of techniques for generating synthetic images from mathematical formulae. The rationale here is that, during the pre-training of vision transformers, feeding such synthetic patterns are sufficient to acquire the necessary visual representations. These images include various types of fractals [1, 17, 19, 20, 28], geometric patterns [18], polygons and other basic shapes [17]. The complexity and smoothness of these shapes can be adjusted along with the brightness, texture, fill-rate, and other factors that affect the rendered image. Labels can be assigned based on the combination of any of these factors, so a labeled dataset of arbitrary quantity can be generated without human intervention. Furthermore, there is close-to-zero risk of generating images with ethical implications, such as societal bias or copyright infringement.

Another major advantage of synthetic datasets is that the quality of images can be improved continuously, unlike natural datasets which can only be enhanced in quantity. Therefore, we anticipate that we could eventually create a synthetic image dataset that surpass the pre-training effect of JFT-300M, by understanding which properties of synthetic images that contribute to pre-training and improving them. Nakashima *et al.* [28] used fractal images to pre-train vision transformers and found that the attention maps tend to focus on the contours (outlines) rather than the textures. Kataoka *et al.* [17] verified the importance of contours by creating a new dataset from well-designed polygons, and exceeded the pre-training effect of ImageNet-21k with this dataset that consists of only contours. These studies indicate that contours are what matter when pre-training vision transformers. However, these studies covered only a limited design space, due to the difficulty of precisely controlling the geometric properties of the contours in each image.

The present study aims to conduct a systematic and thor-

ough investigation of the design space of contour-oriented synthetic images. We systematically investigate the design space of contours by expressing them as a superposition of sinusoidal waves onto ellipses, as shown in Figure 1a. In the same way that a Fourier series can express arbitrary functions, we can express any contour shape with such a superposition of waves onto an ellipse. Such geometrical concepts have appeared in classical physics, *e.g.* de Broglie’s atomic model [6]. Therefore, we name this new dataset “VisualAtom”, and the method to generate the images “visual atomic renderer”. The visual atomic renderer allows us to exhaustively cover the design space of contour-oriented synthetic images, by systematically varying the frequency, amplitude, and phase of each orbit, along with the number of orbits and degree of quantization the orbits. We vary the range of these parameters to generate datasets with different variety of images. We found that variety of contour shapes is a crucial factor for achieving a superior pre-training effect. Our resulting dataset was able to nearly match the pre-training effect of JFT-300M when fine-tuned on ImageNet-1k, while using only 21M images. We summarize the contributions as follows:

Investigative contribution (Figure 1a): We propose a novel methodology based on circular harmonics that allows us to systematically investigate the design space of contour-oriented synthetic datasets. Identifying the optimal range of frequency, amplitude, and quantization of the contours lead to the creation of a novel synthetic dataset VisualAtom with unprecedented pre-training effect on vision transformers.

Experimental contribution (Figure 1b): We show that pre-training ViT-B with VisualAtom can achieve comparable accuracy to pre-training on JFT-300M, when evaluated on ImageNet-1k fine-tuning. Notably, the number of images used to achieve this level of accuracy was approximately 1/14 of JFT-300M. We also show that VisualAtom outperforms existing state-of-the-art FDSL methods.

Ethical contribution: We will release the synthesized image dataset, pre-trained models, and the code to generate the dataset. This will also allow users with limited internet bandwidth to generate the dataset locally. Moreover, the dataset and model will be released publicly as a commercially available license and not limited to educational or academic usage.

2. Related work

2.1. Vision Transformers and Large-scale Datasets

Since the first ViT paper was published by Dosovitskiy *et al.* [10], there have been numerous extensions and improvements to vision transformers. During the same time, there have also been various improvements to SL, SSL, and FDSL. In this section, we will introduce some of the most relevant work in these fields, with the goal of clarifying the

contribution of the present work.

ViT [10] was the first successful attempt to use transformers [35] in computer vision, with only minor modifications made to the original architecture. One of the main findings of their study was that when ViT is pre-trained on larger datasets, such as ImageNet-21k and JFT-300M, the accuracy of downstream-tasks scales with the size of datasets and models. There have also been efforts to use SSL for pre-training vision transformers, e.g. DINO [7], BEiT [3], MAE [12]. However, the highest accuracy for ImageNet-1k fine-tuning is still obtained through vision transformers with supervised pre-training using JFT-3B [42].

The largest datasets, such as JFT-300M/3B/4B [32, 42] and Instagram-3.5B [27] are all excellent candidates for pre-training vision transformers. However, access to these datasets is limited to a selected group of researchers. To address this issue, the LAION-400M/5B dataset [30, 31] has been made public under a permissive license that allows commercial use. LAION-5B contains 5.85 billion CLIP-filtered image-text pairs, which can be used to train large vision+language models. However, as the LAION dataset is also not curated, its use is subject to the same concerns associated with other large datasets (i.e., the existence of societal bias and privacy and copyright concerns) [5]. These ethical concerns associated with large datasets are a serious problem. For example, one of the most widely used datasets in computer vision – ImageNet [9], contains images that violate privacy protection laws and biased samples, can lead to unfairness in models trained on the dataset [39]. Furthermore, public releases of datasets such as 80M Tiny Images [33] have actually been suspended due to discriminatory labels in the dataset [4].

SSL can avoid issues related to discriminatory labels, and also frees us from the tedious task of labeling these huge datasets. However, SSL has not yet achieved the same level of accuracy as SL. Moreover, SSL still relies on natural images and shares the limitations of large datasets e.g. societal bias, privacy and copyright.

2.2. Formula-Driven Supervised Learning

An alternative method for pre-training vision transformers has been proposed, where synthetic images generated from mathematical formulae are used to perform formula-driven supervised learning (FDSL) [19, 28]. Several follow-up studies on FDSL have been published that extend the application of this approach and that increase the quality of the synthetic datasets [1, 14, 17, 18, 38] FDSL has been evolving rapidly, and the quality of images has improved to the point where it now surpasses the level of pre-training that can be achieved using ImageNet-21k [17]. Further, FDSL has considerable potential because it produces a *dynamic* dataset that can be continuously improved, unlike *static* datasets

such as JFT and LAION in which only the quantity of the dataset can be increased. Another advantage of FDSL is that it is entirely free of the common ethical issues associated with natural images, such as societal bias, privacy, and copyright infringement. FDSL also does not require the downloading of an enormous number of images from a host website, because the same synthetic images can be generated locally by anyone that has the formula. These advantages make it a worthwhile endeavor to take the current FDSL that can match ImageNet-21k, and improve its quality to the point where it can match JFT-300M. However, previous FDSL methods have several outstanding issues: 1) difficult to control the variety of the contour outline of the synthetic images, 2) lack of smoothness of the contour, 3) unclear why and how to generate better synthetic images for pre-training. Therefore, an exhaustive investigation regarding the shape of contours in synthetic images is required.

3. Method

Conceptually inspired by de Broglie’s atomic model in the field of quantum physics, we propose the *visual atomic renderer* which can be used to generate synthetic images for learning visual representations in computer vision tasks. Based on the idea that the various object shapes needed in pre-training can be obtained from combinations of smooth wave functions, we present our method in the following in two steps. First, we define the visual atomic renderer that produces 2D visualization of atoms. Second, we present a pre-training method using the VisualAtom, which is synthetic image dataset generated by our proposed the visual atomic renderer. The visual atomic renderer can assign many variations to the contour outlines and control these variations by changing the parameter ranges.

3.1. Visual Atomic Renderer

The visual atomic renderer \mathcal{V} visualizes an atom \mathcal{A} as an image. More specifically, the visual atomic renderer produces and renders a *visual atom* $\mathcal{V}(\mathcal{A}) \subset \mathbb{R}^2$ from an atom \mathcal{A} . Here, we define \mathcal{V} by parameterizing each component of an atom.

Parameterized orbits. We define the orbit O_k by an ellipse with two parameters $(a_k, b_k) \in \mathbb{R}^2$ in a 2D Euclidean space:

$$O_k = \{(x, y)^\top : F_k(x, y) = 0\} \subset \mathbb{R}^2, \quad (1)$$

where F_k is the implicit function of the ellipse given by

$$F_k(x, y) = \frac{x^2}{a_k^2} + \frac{y^2}{b_k^2} - 1. \quad (2)$$

Notably, O_k has the following parametric representation:

$$O_k = \left\{ \begin{pmatrix} a_k \cos \theta \\ b_k \sin \theta \end{pmatrix} : 0 \leq \theta < 2\pi \right\}. \quad (3)$$




Figure 2. **Visual atomic renderer.** (a) De Broglie’s atomic model consists of electron orbits O_k each with a wave function Φ_k and examples of wave functions Φ for the visual atomic renderer with their visualization $V(O, \Phi)$. The orbit O is colored gray and the visualized wave $V(O, \Phi)$ is colored in red. (b) Example of quantization from $V(O, \Phi)$ to $\bar{V}(O, \Phi)$. (c) Visual atom $\mathcal{V}(\mathcal{A})$ consisting of visualized quantized waves $\bar{V}_k = \bar{V}(O_k, \Phi_k; q)$.

where, θ is an auxiliary parameter. To reduce the number of parameters when making a sequence of orbits, we introduce the interval parameter c and recursively determine the parameter values as $(a_{k+1}, b_{k+1}) = (a_k + c, b_k + c)$ given randomly sampled (a_1, b_1) .

Parameterized waves. The simplest definition of the wave function Φ_k can be given by a single sinusoidal wave:

$$\Phi_k(\theta) = \lambda \sin(n\theta), \quad (4)$$

where, $n \in \mathbb{N}$ is a frequency parameter and $\lambda \in \mathbb{R}$ is an amplitude parameter. Note that n needs to be a natural number in order to satisfy the condition $\Phi_k(\theta) = \Phi_k(\theta + 2\pi)$.

To have a sufficient variety of waves, the visual atomic renderer assumes that each orbit has a mixture of two sinusoidal waves and a small noise as follows:

$$\Phi_k(\theta) = \lambda_1 \sin(n_1\theta) + \lambda_2 \sin(n_2\theta) + \eta\epsilon(\theta), \quad (5)$$

where, $n_1, n_2 \in \mathbb{N}$ are frequency parameters, $\lambda_1, \lambda_2 \in \mathbb{R}$ are amplitude parameters, $\eta \in \mathbb{R}$ is a noise-level parameter, and the function $\epsilon(\theta)$ returns a one-dimensional Perlin noise that varies randomly for each θ . Examples of wave functions (w/o noise) are shown in Figure 2a. This simple definition using sinusoidal waves will provide diverse shapes that help ViT to learn visual representations.

Visual waves. Given an orbit O_k with a wave function Φ_k , we define the visual wave $V(O_k, \Phi_k) \subset \mathbb{R}^2$ by multiplying the wave function by the parametric representation of the orbit in Eq. (3) as follows:

$$V(O_k, \Phi_k) = \{\Phi_k^*(\theta) : 0 \leq \theta < 2\pi\}, \quad (6)$$

where,

$$\Phi_k^*(\theta) = \begin{pmatrix} a_k(\Phi_k(\theta) + 1) \cos \theta \\ b_k(\Phi_k(\theta) + 1) \sin \theta \end{pmatrix} \in \mathbb{R}^2. \quad (7)$$

Examples of visual waves are shown in Figure 2a. Interestingly, we see that some of them look like the shapes of natural objects, such as a star, a flower, a bird, and an amoeba.

Further, to incorporate sharp corners, we introduce a quantization parameter q , which quantizes $\Theta = [0, 2\pi)$ into a discrete set of points $\bar{\Theta} = \{\theta_i : i = 0, 1, 2, \dots, q-1\}$, where $\theta_i = 2\pi i/q$. We define the quantized visual wave:

$$\bar{V}(O_k, \Phi_k; q) = \bigcup_{i=0}^{q-1} \ell(\mathbf{x}_i, \mathbf{x}_{i+1}), \quad (8)$$

where, $\mathbf{x}_i = \Phi_k^*(\theta_i)$ is the i -th point on the visual wave and

$$\ell(\mathbf{x}_i, \mathbf{x}_{i+1}) = \{\alpha\mathbf{x}_i + (1-\alpha)\mathbf{x}_{i+1} : 0 \leq \alpha \leq 1\} \quad (9)$$

is the line between the two points (\mathbf{x}_i and \mathbf{x}_{i+1}). This procedure is illustrated in Figure 2b.

Finally, given an atom \mathcal{A} , we define a visual atom $\mathcal{V}(\mathcal{A})$ by taking the union set of quantized visual waves as follows:

$$\mathcal{V}(\mathcal{A}) = \bigcup_{k=1}^K \bar{V}(O_k, \Phi_k; q). \quad (10)$$

An example of the visual atom is shown in Figure 2c.

Rendering. From Eqs. (8) and (10), a visual atom $\mathcal{V}(\mathcal{A})$ can be viewed as a set of lines. Thus, we render the lines to synthesize the image. For rendering, there are two parameters: line thickness l_t and line color $l_c \in [0.0, 1.0]$ in gray scale. The line thickness is fixed to $l_t = 1$ pixel as the baseline. The color is randomly selected for each orbit. Finally, the nucleus (the center of the visual atom) is parameterized by a 2D vector $\mathbf{p} \in \mathbb{R}^2$ by which the rendering position is determined. All of the lines are rendered on a black background.

Compiling VisualAtom. We create a pre-training dataset D , which consists of pairs $(I, y) \in D$ of an image I and a class label $y \in \{1, 2, \dots, C\}$, by the following two steps.

Table 1. Parameters of the visual atomic renderer. Gr: group id by which the parameters are divided into two groups.

Gr	Parameter	Baseline Range
1	Number of orbits K	$\{1, 2, \dots, 200\}$
	Orbit parameters a_1, b_1	$[1.0, 400.0]$
	Orbit interval parameter c	1.0 (fixed)
	Frequency parameters n_i	$\{0, 1, \dots, 20\}$
	Amplitude parameters λ_i	0.5 (fixed)
	Quantization parameter q	$\{200, \dots, 1000\}$
	Noise level η	$[0.0, 1.0]$
2	Line-thickness parameter l_t	1 (fixed)
	Line-color parameter l_c	$[0.0, 1.0]$
	Nucleus position \mathbf{p}	$[-1.0, 1.0]^2$

First, C independent sets of parameters are randomly sampled to determine the classes of the visual atoms. Here, we divide the parameters into two groups, as shown in Table 1. For each class, the parameters in the first group are randomly sampled and fixed. Second, N images are generated for each class. Here, the parameters of the second group are randomly selected for each image, as this increases the intra-class variability. Finally, the training dataset consists of $|D| = NC$ images. We name this synthesized image dataset “VisualAtom”. The cross-entropy loss over D is minimized during the pre-training phase.

4. Experiments

We conduct experiments to validate the effectiveness of VisualAtom for pre-training vision transformers. We first explore the relationship between various contour features and the pre-training effect, by changing the parameters of the visual atomic renderer such as frequency in Section 4.1.

We then examine the fine-tuning performance in comparison with state-of-the-art pre-training methods, including supervised and self-supervised pre-training on ImageNet-21k and JFT-300M in Section 4.2. Finally, we discuss considerations and limitations of our method in Section 4.3.


4.1. Pre-training with VisualAtom

We first validate the effectiveness of VisualAtom by comparing with ExFractalDB/RCDB [17]. Here, both ExFractalDB-1k/RCDB-1k are state-of-the-art datasets for FDSL that each comprise one million images of well-designed fractals/polygons. For a fair comparison with ExFractalDB-1k/RCDB-1k in terms of the number of images, we compiled VisualAtom-1k, a dataset comprised of one million images of visual atoms, where the number of classes is $C = 1,000$ and the number of images per class is $N = 1,000$. We also report results when using SSL on ImageNet-1k [9]. Pre-training and fine-tuning were done with almost the same hyper-parameters as DeiT [34].

Table 2 compares the accuracy, using the following four

Table 2. Comparison of fine-tuning accuracy on four datasets: CIFAR10 (C10), CIFAR100 (C100), ImageNet-100 (IN100) and ImageNet-1k (IN-1k). Cross entropy loss is used for SL. DINO [7] is used for SSL. ViT-Tiny with a patch size of 16 is used for all experiments. These results are also included in Table 7. Best scores in FDSL are shown in bold.

Pre-training	C10	C100	IN100	IN-1k
Scratch	78.3	57.7	73.2	72.6
ImageNet-1k (SSL)	97.7	82.4	89.0	-
ExFractalDB-1k	97.2	81.8	88.1	73.0
RCDB-1k	97.0	82.2	88.5	73.1
VisualAtom-1k (ours)	97.6	84.9	90.3	74.2



datasets: CIFAR10 (C10) [22], CIFAR100 (C100) [22], ImageNet-100 (IN100)¹, and ImageNet-1k (IN-1k). The findings showed that the proposed VisualAtom-1k outperforms RCDB-1k on all of the four fine-tuning datasets. Notably, we obtained 2.7 point improvement on the C100 dataset (VisualAtom-1k 84.9 v.s. RCDB-1k 82.2). In addition, the performance obtained with VisualAtom-1k was comparable to, or better than, SSL (DINO). Overall, the results confirmed that VisualAtom-1k achieves state-of-the-art performance when used with FDSL.

We then conducted experiments to analyze the effects of individual parameters of the visual atomic renderer, and discuss the importance of each contour feature. The baseline ranges from which parameter values were sampled are summarized in Table 1.

Frequency Parameters (Table 3). The two frequency parameters n_1 and n_2 in Eq. (5) are the most important parameters because they determine the overall shape of visual atoms. Table 3 shows the results of experiments conducted using various frequency ranges. Because n_1 and n_2 need to be natural numbers, we varied the minimum and maximum numbers for them. For example, the setting $[\min, \max] = [0, 20]$ at the first row indicates that n_1 and n_2 are randomly sampled from $\{0, 1, \dots, 20\}$.

The results show that the best setting is $[0, 20]$ and that VisualAtom-1k outperforms RCDB-1k at all settings. Increasing the maximum number in the range did not improve the pre-training effect because the frequency of the sinusoidal waves became too high. The value of 20 was considered to be a reasonable maximum frequency for pre-training with an image size of 224×224 . In addition, the fine-tuning accuracy was tested when the maximum frequency is fixed at 20, and the minimum frequency is increased grad-

¹This is a subset of ImageNet with 100 object categories.

Table 3. Accuracy when varying the range of frequency parameters n_1, n_2 .

Range of n, m		C10	C100	IN100
min	max			
0	20	97.6	84.9	90.3
0	40	97.1	84.3	89.5
0	60	97.3	84.1	89.1
2	20	97.6	84.8	90.3
10	20	97.5	84.4	89.8
20	20	97.3	83.6	89.6




Table 4. Accuracy when varying the range of number of orbits K .

Range of K		C10	C100	IN100
min	max			
1	200	97.6	84.9	90.3
20	200	97.5	84.7	89.9
100	200	97.5	84.5	89.8
200	200	97.5	84.3	89.4





Table 5. Accuracy when varying the range of quantization parameter q .

Range of q		C10	C100	IN100
min	max			
200	1,000	97.6	84.9	90.3
800	1,000	97.4	85.1	89.9
3	200	97.3	84.6	89.7
3	500	97.3	84.9	90.1
3	1,000	97.4	85.0	90.1



ually from 0 to 20 to narrow the frequency range. The results showed that the fine-tuning accuracy declined for all downstream-tasks as the frequency range was narrowed. We speculate that this decline in accuracy may be caused by the decrease in category variations determined by the combination of the two frequency parameters.

Number of Orbits Parameter (Table 4). The parameter K defines the number of orbits on one-image. K is sampled uniformly per category from pre-defined natural number range or is fixed at a pre-defined natural number. Table 4 shows the results when varying the range of K .

Using a maximum K of 200, we found that the fine-tuning accuracy tends to increase for almost all downstream-tasks at the lower minimum K and the larger variation of K per category. This finding supports the hypothesis that the pre-training dataset should contain categories of various size object shapes.

Quantization Parameter (Table 5). The quantization parameter q controls the smoothness of the contours. As shown in Table 5, we did not observe any degradation of accuracy even for the smallest range of [800, 1000]. On the other hand, when q is in the range of [3, 200], we observed a drop in accuracy. Furthermore, we observed that the accuracy improved as the minimum q was fixed to 1 and the maximum q was increased to 500 or 1000.


These results indicate that the quantization parameter q needs to be above a certain value, for better representation and smoothness of the contours.

Amplitude Parameters (Table 6). The two amplitude parameters, λ_1 and λ_2 define the amplitude of two separate sinusoidal waves added on the orbit. λ_1 and λ_2 are sampled uniformly per category from a pre-defined non-negative decimal range or fixed at a pre-defined decimal. Table 6 shows the results of experiments with various amplitude ranges.

The findings showed that varying the amplitude for each category and separating the distance between categories had no effect on the pre-training performance. Therefore, for

Table 6. Accuracy when varying the range of amplitude parameters λ_1, λ_2 .

Range of λ_1, λ_2		C10	C100	IN100
min	max			
0.5	0.5	97.6	84.9	90.3
0.0	0.5	97.4	84.9	89.5
1.0	1.0	97.3	84.4	89.6
0.5	1.0	97.4	84.5	89.6
0.0	1.0	97.4	84.4	89.8



simplicity, we fixed the amplitude at 0.5 for the baseline configuration of VisualAtom.

Best Practice in VisualAtom Pre-training. According to the results of the experiments above plus some additional experiments not shown here for tuning the parameter combinations, the highest accuracy occurred in category (1000/21,000), instance (1,000), number of orbits (1, 2, ..., 200), vertical and horizontal diameters of orbits ([1.0, 400.0]), orbits spacing (1.0), frequency (0, 1, ..., 20), amplitude (0.5), quantization (200, 201, ..., 1000), noise level ([0.0, 1.0]), line-thickness (1), line-color ([0.0, 1.0]), nucleus position ([-1.0, 1.0]²), image size (512 × 512) and training epochs (300/90). This best combination of parameters (baseline parameters) is shown in Table 1. The fine-tuning accuracy is shown in Table 2 and Table 7.

4.2. Comparisons

For comparison with previous works, we use seven datasets for evaluation: Stanford Cars (Cars) [21], Flowers [29], Pascal VOC 2012 (VOC12) [11], Places 30 (P30) [19], C10 [22], C100 [22] and IN100 [19]. To compare the performance of ViT-B pre-trained on JFT-300M in the original ViT paper [10], we also report results of ViT-B for the fine-tuning performed using ImageNet-1k [9] with resolution of 384 × 384.

Table 7. Comparison of pre-training methods. Best and second-best scores at ViT-Tiny are shown in underlined bold and bold, respectively. In comparison at ViT-Base, higher scores are shown in bold.

Model	Pre-training	Type	C10	C100	Cars	Flowers	VOC12	P30	IN100	Average
ViT-T	Scratch	–	78.3	57.7	11.6	77.1	64.8	75.7	73.2	62.6
	ImageNet-1k	SL	98.0	85.5	89.9	99.4	88.7	80.0	–	–
	ImageNet-1k	SSL (DINO)	97.7	82.4	88.0	98.5	74.7	78.4	89.0	86.9
	PASS	SSL (DINO)	97.5	84.0	86.4	98.6	82.9	79.0	82.9	87.8
	FractalDB-1k [28]	FDSL	96.8	81.6	86.0	98.3	80.6	78.4	88.3	87.1
	ExFractalDB-1k	FDSL	97.2	81.8	87.0	98.9	80.6	78.0	88.1	87.4
	RCDB-1k	FDSL	97.0	82.2	86.5	98.9	80.9	79.7	88.5	87.6
	VisualAtom-1k (ours)	FDSL	97.6	84.9	88.8	98.9	82.0	81.2	90.3	89.1
ViT-B	RCDB-21k	FDSL	96.8	82.9	85.9	99.0	81.2	81.2	90.2	88.2
	VisualAtom-21k (ours)	FDSL	97.7	86.7	89.2	99.0	82.4	81.6	91.3	89.7

Table 8. Comparison with fine-tuning accuracy on ImageNet-1k. Res. indicates the image resolution at fine-tuning. Best and second-best scores at each resolution and model scale are shown in underlined bold and bold, respectively.

Pre-training	Res.	Type	ViT-T	ViT-B
Scratch	224 ²	–	72.6	79.8
ImageNet-21k	224 ²	SL	74.1	81.8
ExFractalDB-1k	224 ²	FDSL	73.7	80.4
ExFractalDB-21k	224 ²	FDSL	73.6	82.7
RCDB-1k	224 ²	FDSL	73.1	82.3
RCDB-21k	224 ²	FDSL	72.8	82.4
VisualAtom-1k (ours)	224 ²	FDSL	74.2	82.3
VisualAtom-21k (ours)	224 ²	FDSL	73.8	82.7
ImageNet-21k	384 ²	SL	–	83.0
JFT-300M [10]*	384 ²	SL	–	84.2
VisualAtom-21k (ours)	384 ²	FDSL	–	83.7

* Since the JFT-300M, which is not available, cannot be compared with the same setting, we transcribe the accuracy with the different setting reported in the previous work [10].

Table 9. Fine-tuning results on COCO dataset. Best values for each learning type are shown in bold.

Pre-training	COCO Det			COCO Inst Seg		
	AP ₅₀	AP	AP ₇₅	AP ₅₀	AP	AP ₇₅
Scratch	63.7	42.2	46.1	60.7	38.5	41.3
ImageNet-1k	69.2	48.2	53.0	66.6	43.1	46.5
ImageNet-21k	70.7	48.8	53.2	67.7	43.6	47.0
ExFractalDB-1k	69.1	48.0	52.8	66.3	42.8	45.9
ExFractalDB-21k	69.2	48.0	52.6	66.4	42.8	46.1
RCDB-1k	68.3	47.4	51.9	65.7	42.2	45.5
RCDB-21k	67.7	46.6	51.2	64.8	41.6	44.7
VisualAtom-1k	68.8	47.8	52.4	66.0	42.6	45.8
VisualAtom-21k	68.3	47.4	52.3	65.3	42.2	45.4

Comparison with Previous Works (Table 7). In Table 7, we show the fine-tuning accuracies obtained using the seven datasets. Here, the following two experiments are conducted: a one-million image-scale experiment with ViT-T and a ten-million image-scale experiment with ViT-B.

In the **one-million image-scale experiment with ViT-T**, the following pre-training methods were compared: 1) SL with ImageNet-1k, 2) SSL (DINO [7]) with ImageNet-1k/PASS [2], 3) FDSL with FractalDB-1k [19], ExFractalDB-1k [17], RCDB-1k [17], and VisualAtom-1k.

The results in Table 7 show that VisualAtom-1k achieved the highest average score² for the SSL and FDSL methods. We observed a 1.3 points performance gap between FDSL with VisualAtom-1k (89.1) and SSL with PASS (87.8). This result is surprising given that VisualAtom-1k contains 1.0M synthetic images and PASS contains 1.4M real images. Moreover, all of the performance rates obtained with VisualAtom-1k were either equal to, or better than, those obtained with the other FDSL datasets. The finding showed that FDSL with VisualAtom-1k partially outperforms SL with ImageNet-1k on P30 (81.2 vs. 80.0).

In the **ten-million image-scale experiments with ViT-B**, the two FDSL datasets of RCDB and VisualAtom were scaled up to 21,000 categories. The findings showed that the average accuracy is further improved from 89.1 (VisualAtom-1k) to 89.7 (VisualAtom-21k), and that the score obtained using VisualAtom-21k was consistently equal to, or better than, that obtained with RCDB-21k.

ImageNet-1k fine-tuning (Table 8). Table 8 shows fine-tuning results obtained using ImageNet-1k. With ViT-T, we see that VisualAtom-1k surpassed ImageNet-21k (74.2 v.s. 74.1). We also see that relatively small-size datasets are effective for effectively pre-training ViT-T with FDSL. With ViT-B, we see that the scaled-up VisualAtom-21k improved the accuracy from 82.3 to 82.7, which is comparable to that

²The average accuracy for SL pre-training was not calculated in the previous paper [17].

of ExFractalDB-21k and is better than that of ImageNet-21k. Finally, we conducted an experiment using pixel resolution of 384×384 in order to compare results with JFT-300M, a massive milestone research dataset that was originally used to pre-train vision transformers. Although the VisualAtom-21k did not outperform JFT-300M, it achieved a notable performance with 14.2 times fewer synthetic images (VisualAtom-21k 83.7 with 21M synthetic images vs. JFT-300M 84.2 with 300M real images).

COCO detection/instance segmentation (Table 9). We also conducted validation experiments with object detection and instance segmentation using the COCO [24] dataset. We used Swin Transformer (Base) [26] for the backbone model and Mask R-CNN [13] for the head. We pre-trained Swin-Base on VisualAtom-1k/21k for 300/90 epochs and performed fine-tuning using the COCO dataset for 60 epochs. Table 9 shows a comparison of our VisualAtom with previous SL and FDSL methods, based on AP scores for detection and instance segmentation tasks. In the fine-tuning with detection and instance segmentation tasks, the results showed that VisualAtom has a higher pre-training effect than RCDB and slightly lower pre-training effect than ExFractalDB. This order was also observed when the dataset category size were scaled up to 21k categories.

4.3. Discussion and limitations

What should we focus on in FDSL? In Section 4.1, we conducted systematic experiments to investigate the relationship between various contour features and the pre-training effect, by changing the parameters of the visual atomic renderer to control contour features of visual atoms. We found that extending the variation of the following parameters enhances the pre-training effect; frequency (controlling the shape of contours), number of orbits (controlling the size of contours) and quantization (controlling the smoothness of contours). In particular, increasing the variation of frequency, in other words, using various contour shapes, greatly improved the pre-training effect. With the insight we gained from this systematic investigation, we are now able to construct better FDSL datasets. The ability to continuously improve the quality of the images and not only the quantity is what distinguishes FDSL from SL with real images.

Potential of VisualAtom as a FDSL dataset. In Section 4.2, we validated the pre-training effect of FDSL using our proposed VisualAtom by comparing with many other benchmarks. We found that VisualAtom enhances the pre-training effect for relatively small datasets. In particular, fine-tuning comparisons of ViT-T pre-trained on VisualAtom-1k outperformed all other previous FDSL datasets, including ExFractalDB-1k and RCDB-1k (e.g., VisualAtom-1k 89.1 vs. RCDB-1k 87.6, with average accuracies shown in Table 7). The ViT-T pre-trained

on VisualAtom-1k is better than ViT-T pre-trained on ImageNet-21k in Table 8. Although a balance exists between model size and pre-training dataset size, the proposed method is also more accurate than ViT-T pre-trained using RCDB-1k/ExFractalDB-1k. Moreover, when conditions were aligned with ViT-B pre-trained using JFT-300M, our proposed VisualAtom-21k differed by only a 0.5 point gap on ImageNet-1k fine-tuning with 1/14 the number of images. These findings indicate that approaches using contour-based FDSL datasets have considerable potential to enhance the pre-training effect of ViT.

Limitation. VisualAtom can still be improved. For example, even though one-million image-scale pre-training significantly improved the recognition performance with ViT-T and relatively small down-stream tasks (see Table 7 showing the results obtained using our VisualAtom-1k with ViT-T), ten-million image-scale pre-training cannot surpass the related FDSL dataset in Table 8. At this stage, we speculate that fractal-like, recursively intricate patterns could conceivably increase the complexity of visual atoms and bring about further improvements. We intend to further explore ways in which complex patterns can be generated efficiently in order to extend the effectiveness of pre-training without using any real images or human supervision.

5. Conclusion

In this study, we proposed the visual atomic renderer and how it can control contour features of synthesized images from a mixture of sinusoidal waves. The simplicity and flexibility of the visual atomic renderer allowed us to perform a systematic exploration of the design space of contour-oriented synthetic datasets. We found that the rich variation of contour shapes among categories enhances the pre-training effect. We also found that our proposed VisualAtom dataset outperforms all existing FDSL datasets. In a fine-tuning experiment using ImageNet-1k, FDSL with VisualAtom-21k achieved a top-1 accuracy of 83.7%, and closed the gap in accuracy between FDSL and SL on JFT-300M to 0.5%. In addition, compared to JFT-300M, VisualAtom-21k, has only 1/14 the number of images.

Acknowledgement

This paper is based on results obtained from a project, JPNP20006, commissioned by the New Energy and Industrial Technology Development Organization (NEDO). Computational resource of AI Bridging Cloud Infrastructure (ABCI) provided by National Institute of Advanced Industrial Science and Technology (AIST) was used. We want to thank Junichi Tsujii, Yutaka Satoh, Ryo Nakamura, Ryosuke Yamada, Kodai Nakashima, Xinyu Zhang, Zhaoqing Wang, Toshiki Omi, Seitaro Shinagawa, Koshi Maki-hara for their helpful comments in research discussions.

References

- [1] Connor Anderson and Ryan Farrell. Improving Fractal Pre-training. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1300–1309, 2022. [2](#), [3](#)
- [2] Yuki M Asano, Christian Rupprecht, Andrew Zisserman, and Andrea Vedaldi. PASS: An imagenet replacement for self-supervised pretraining without humans. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. [7](#)
- [3] Hangbo Bao, Li Dong, and Furu Wei. BEiT: Bert Pre-training of Image Transformers. *arXiv preprint arXiv:2106.08254*, 2021. [3](#)
- [4] Abeba Birhane and Vinay Uday Prabhu. Large Image Datasets: A Pyrrhic Win for Computer Vision? In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1536–1546. IEEE, 2021. [3](#)
- [5] Abeba Birhane, Vinay Uday Prabhu, and Emmanuel Kahembwe. Multimodal Datasets: Misogyny, Pornography, and Malignant Stereotypes. *arXiv preprint arXiv:2110.01963*, 2021. [2](#), [3](#)
- [6] Louis de Broglie. *Philosophical Magazine*, 47:446–458, 1924. [2](#)
- [7] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9650–9660, 2021. [3](#), [5](#), [7](#)
- [8] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision Transformer Adapter for Dense Predictions. *arXiv preprint arXiv:2205.08534*, 2022. [1](#)
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [3](#), [5](#), [6](#)
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [11] Mark Everingham, SM Eslami, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The Pascal Visual Object Classes Challenge: A Retrospective. *International journal of computer vision*, 111(1):98–136, 2015. [6](#)
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked Autoencoders are Scalable Vision Learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. [3](#)
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [8](#)
- [14] Nakamasa Inoue, Eisuke Yamagata, and Hirokatsu Kataoka. Initialization Using Perlin Noise for Training Networks with a Limited Amount of Data. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 1023–1028. IEEE, 2021. [3](#)
- [15] Jitesh Jain, Anukriti Singh, Nikita Orlov, Zilong Huang, Jiachen Li, Steven Walton, and Humphrey Shi. SeMask: Semantically Masked Transformers for Semantic Segmentation. *arXiv preprint arXiv:2112.12782*, 2021. [1](#)
- [16] Longlong Jing and Yingli Tian. Self-supervised Visual Feature Learning With Deep Neural Networks: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(11):4037–4058, 2021. [2](#)
- [17] Hirokatsu Kataoka, Ryo Hayamizu, Ryosuke Yamada, Kodai Nakashima, Sora Takashima, Xinyu Zhang, Edgar Josafat Martinez-Noriega, Nakamasa Inoue, and Rio Yokota. Replacing Labeled Real-Image Datasets With Auto-Generated Contours. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21232–21241, 2022. [2](#), [3](#), [5](#), [7](#)
- [18] Hirokatsu Kataoka, Asato Matsumoto, Ryosuke Yamada, Yutaka Satoh, Eisuke Yamagata, and Nakamasa Inoue. Formula-driven Supervised Learning with Recursive Tiling Patterns. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4098–4105, 2021. [2](#), [3](#)
- [19] Hirokatsu Kataoka, Kazushige Okayasu, Asato Matsumoto, Eisuke Yamagata, Ryosuke Yamada, Nakamasa Inoue, Akio Nakamura, and Yutaka Satoh. Pre-training without Natural Images. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, November 2020. [2](#), [3](#), [6](#), [7](#)
- [20] Hirokatsu Kataoka, Kazushige Okayasu, Asato Matsumoto, Eisuke Yamagata, Ryosuke Yamada, Nakamasa Inoue, Akio Nakamura, and Yutaka Satoh. Pre-training without Natural Images. *International Journal of Computer Vision (IJCV)*, 130(2):990–1007, 2022. [2](#)
- [21] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3D Object Representations for Fine-grained Categorization. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 554–561, 2013. [6](#)
- [22] Alex Krizhevsky, Hinton, and Geoffrey. Learning Multiple Layers of Features from Tiny Images. 2009. [5](#), [6](#)
- [23] Feng Li, Hao Zhang, Huaizhe Xu, Shilong Liu, Lei Zhang, Lionel M. Ni, and Heung-Yeung Shum. Mask DINO: Towards A Unified Transformer-based Framework for Object Detection and Segmentation. *arXiv preprint arXiv:2206.02777*, 2022. [1](#)
- [24] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. [8](#)
- [25] Ze Liu, Han Hu, Yutong Lin, Zhenda Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer V2: Scaling Up Capacity and Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12009–12019, 2022. [1](#)

- [26] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical Vision Transformer using Shifted Windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 8
- [27] Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the Limits of Weakly Supervised Pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pages 181–196, 2018. 2, 3
- [28] Kodai Nakashima, Hirokatsu Kataoka, Asato Matsumoto, Kenji Iwata, Nakamasa Inoue, and Yutaka Satoh. Can Vision Transformers Learn without Natural Images? In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 1990–1998, 2022. 2, 3, 7
- [29] Maria-Elena Nilsback and Andrew Zisserman. Automated Flower Classification over a Large Number of Classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008. 6
- [30] Christoph Schuhmann, Romain Beaumont, Cade W Gordon, Ross Wightman, mehdi cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Mitchell Wortsman, Richard Vencu, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: An Open Large-Scale Dataset for Training Next Generation Image-Text Models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 2, 3
- [31] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. LAION-400M: Open Dataset of CLIP-Filtered 400 Million Image-Text Pairs. *arXiv preprint arXiv:2111.02114*, 2021. 3
- [32] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting Unreasonable Effectiveness of Data in Deep Learning Era. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 1, 2, 3
- [33] Antonio Torralba, Rob Fergus, and William T Freeman. 80 million tiny images: A large data set for nonparametric object and scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 30(11):1958–1970, 2008. 3
- [34] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training Data-efficient Image Transformers & Distillation through Attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021. 5
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All You Need. *Advances in neural information processing systems*, 30, 2017. 3
- [36] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, and Furu Wei. Image as a Foreign Language: BEiT Pretraining for All Vision and Vision-Language Tasks. *arXiv preprint arXiv:2208.10442*, 2022. 1
- [37] Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Carmon Yair, Simon Kornblith, and Ludwig Schmidt. Model soups: Averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. *arXiv preprint arXiv:2203.05482*, 2022. 1
- [38] Ryosuke Yamada, Hirokatsu Kataoka, Naoya Chiba, Yukiyasu Domae, and Tetsuya Ogata. Point Cloud Pre-Training With Natural 3D Structures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21283–21293, 2022. 3
- [39] Kaiyu Yang, Klint Qinami, Li Fei-Fei, Jia Deng, and Olga Russakovsky. Towards Fairer Datasets: Filtering and Balancing the Distribution of the People Subtree in the ImageNet Hierarchy. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 547–558, 2020. 2, 3
- [40] Kaiyu Yang, Jacqueline Yau, Li Fei-Fei, Jia Deng, and Olga Russakovsky. A Study of Face Obfuscation in ImageNet. In *International Conference on Machine Learning (ICML)*, 2022. 2
- [41] Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. Coca: Contrastive Captioners are Image-Text Foundation Models. *arXiv preprint arXiv:2205.01917*, 2022. 1
- [42] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling Vision Transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12104–12113, 2022. 1, 2, 3