
ON HATE SCALING LAWS FOR DATA-SWAMPS

Abeba Birhane*

Mozilla Foundation, San Francisco, USA &
School of Computer Science and Statistics
Trinity College Dublin, Ireland
birhanea@tcd.ie

Vinay Prabhu*, Sang Han

Independent researchers
San Francisco
USA

vinaypra@alumni.cmu.edu, sanghan@protonmail.com

Vishnu Naresh Boddeti

Computer Science and Engineering
Michigan State University
vishnu@msu.edu

ABSTRACT

‘Scale the model, scale the data, scale the GPU-farms’ is the reigning sentiment in the world of generative AI today. While model scaling has been extensively studied, data scaling and its downstream impacts remain under explored. This is especially of critical importance in the context of visio-linguistic datasets whose main source is the World Wide Web, condensed and packaged as the *CommonCrawl* dump. This large scale data-dump, which is known to have numerous drawbacks, is repeatedly mined and serves as the data-motherlode for large generative models. In this paper, we: 1) investigate the effect of scaling datasets on hateful content through a comparative audit of the LAION-400M and LAION-2B-en, containing 400 million and 2 billion samples respectively, and 2) evaluate the downstream impact of scale on visio-linguistic models trained on these dataset variants by measuring racial bias of the models trained on them using the Chicago Face Dataset (CFD) as a probe. Our results show that 1) the presence of hateful content in datasets, when measured with a Hate Content Rate (HCR) metric on the inferences of the Pysentimiento hate-detection Natural Language Processing (NLP) model, *increased by nearly 12%* and 2) societal biases and negative stereotypes were also exacerbated with scale on the models we evaluated. As scale increased, the tendency of the model to associate images of human faces with the ‘human being’ class over 7 other offensive classes *reduced by half*. Furthermore, for the Black female category, the tendency of the model to associate their faces with the ‘criminal’ class *doubled*, while *quintupling* for Black male faces. We present a qualitative and historical analysis of the model audit results, reflect on our findings and its implications for dataset curation practice, and close with a summary of our findings and potential future work to be done in this area. All the meta-datasets curated in this endeavor and the code used are shared at: https://github.com/vinayprabhu/hate_scaling.

Content warning: This article contains examples of hateful text and NSFW images that might be disturbing, distressing, and/or offensive.

1 Introduction

Generative AI models have come to captivate diverse stakeholders, spanning from researchers [13, 60, 53], to media institutions [41, 45], and even large-scale investment firms [27, 50]. This trend can be traced back to the emergence of Dall.E [80], a text-to-image visio-linguistic model released in April 2022, which purportedly attracted over a million users within the first three months of its launch, and was celebrated with claims like: “[t]he first AI technology that has caught fire with regular people” [41]. Subsequently, models such as StableDiffusion [81] and Midjourney emerged, followed by black box projects from Big Tech such as Imagen [83], Parti [108], and BASIC [74], access to which was never given to the general public. While Stable Diffusion and its variants have been trained on the open-sourced

*Equal contribution

datasets from the LAION family, little is known about the datasets that were used to train models such as Dall-E, Parti [108], and Imagen [83].

Fundamental to this multimodal model boom is large-scale visio-linguistic datasets containing image-text pairs, which form the main focal point of this paper. Broadly speaking, these datasets are of two types: those that are open-source, “freely available” and mainly scraped from the Common Crawl (such as LAION-400M [88] and LAION-5B [87]), and those that are closed datasets curated internally by Big Tech corporate labs (such as Google’s ALIGN 1.7B/ALIGN 6.6B [49], JFT-5B [74], and OpenAI’s WebImageText-WIT [78]). The latter remain outside the reach of independent audits and evaluations, while the models trained on such datasets are public-facing and commercialized via APIs (such as Microsoft Bing image-creator powered by Dall.E, or the Dall.E API). These models are also being adopted in various commercial tools and applications such as stock photo generation [57], which contribute to accelerating their adoption and usage.

The open-source variants of these datasets are getting bigger and now breaching the billion-samples mark for two reasons: firstly, there is the unquestioned subservience to the *scale is all you need* mandate handed down from Big Tech disseminations [49, 13] that forms the motivational drive. Secondly, there is the emergent nexus between dataset curation and venture capital resulting in capital infusion into these dataset curation efforts which was hitherto missing: the LAION-5B [87] dataset, for example, was sponsored by Hugging Face, Doodlebot and Stability.ai, as per their blog-post announcement.

In turn, this *scale is all you need* mandate emerges from two schools of reasoning in published literature venerating scale. The first pertains to vague “dataset scaling laws” that we cover in detail in Appendix B and the second pertains to the non-reproducible empirical results buried in subsections of some of the canonical papers which describe the (closed) datasets used for training models such as ALIGN [49], Imagen [83] and Parti [108] (See Appendix C). We also note that these high-profile disseminations are increasingly turning into a flag-posting exercise that involves tactfully concealing critical details on the manner in which the dataset was curated and where the data came from (See Appendix D for a deeper exploration).

All of this has resulted in a “*scrape-first-ask-questions-later*” data creation and curation culture, generating plausibly illegal gargantuan datasets and models, that has in turn elicited a slew of copyright lawsuits [54], en masse fetishization of women’s bodies in an emergent synthetic digital culture [70], outright bans of model outputs from art-forums [25], and marquee datasets filled with duplicates [105]. These dataset and model concerns and drawbacks, in turn, result in downstream negative impacts, often against marginalized groups, for example: exacerbation of negative stereotypes and biases [59, 31, 13], discriminatory and harmful representation [97, 7, 106] and cultural and linguistic homogeneity [6, 8, 15].

Hateful, abusive, racist, aggressive and targeted speech are overlapping phenomena yet each can be characterized along dimensions such as directed, generalized, explicit and implicit abuse [104]. Oftentimes, the common targets of hateful speech are minoritized groups. Based on analysis of generated data to improve hate detection, [100] highlight that most of the common targets of hate include Black people, women, Muslims and trans people. Furthermore, hateful, abusive, and aggressive speech is a systemic problem. [22], for example, examined hate speech and abusive language detection datasets and found systematic racial bias in all datasets. Subsequently, they found that classifiers trained on them predict tweets written in African-American English as abusive at a substantially higher rate. Similarly, [1] studied the outputs generated from GPT-3 when the word “Muslim” is included in the prompt. They found that 66 out of the 100 completions are violent, where these violent completions are less likely for other religions. While issues concerning hate speech are rooted in systemic structures, current attention has focused too much on finding toxic language, performance maximization and engineering solutions. [76], thus argue more attention ought to be paid to recognizing the root causes and focus on the social and ethical initiatives along with the real-world impacts of hate-speech.

In this paper, we examine: 1) the impact of scale on hate-speech through audits of textual descriptions in two datasets: LAION-400M and LAION-2B-en, and 2) the downstream negative impact on models trained on these two datasets through audits of such models trained on their variants.

The rest of the paper is organized as follows. In Section 2, we survey previous work on scale within the broader technical landscape as well as within the Machine Learning (ML) community. In Section 3, we present our dataset audit methodology followed by our findings in Section 4 which reveal that hateful content increased when the dataset size was scaled from 400 million to 2 billion. To establish this, we used an NLP-aided quality audit of the datasets by measuring the hate content in the alt-text image descriptions using a state-of-the-art (SoTA) pre-trained open-source model named *Pysentimiento* [72]. We then focus on the downstream consequences of the hate-scaling phenomenon by measuring the racial biases exhibited by visio-linguistic models that have been trained on these two datasets, where we detail our audit methodology and experimental design in Section 5. Our findings, summarized

in Section 6, demonstrate that associations of Black people’s faces with dehumanizing classes such as criminals markedly increased when the model size and architecture are held constant and the dataset is scaled from 400 million to 2 billion. We delve deeper and position our experimental findings in historical context by qualitatively examining historical patterns of dehumanization and criminalization of Black bodies in Section 7. Section 8 provides some caveats, discussions, and recommendations for equitable, responsible and accountable dataset creation, curation, and management practice. We then highlight a number of extensions for future work and conclude in Section 9.

2 Scale: An Overview

Current thinking around scale can be broadly categorized under two differing approaches: that which sees scale as a solution to problems such as model performance and generalization, and that which emphasizes numerous concerns that arise with an unwavering commitment to scale. We present both below.

2.1 Scale as a Solution

The race to scale is a fixation driving not only research in ML but also the larger tech “innovation” discourse. Entrepreneurs are warned that “*if you don’t know how to scale, don’t innovate*” [73, 90]. Marked by the taken for granted, field-wide concept of *scaling laws*, large scale is thought to correlate with better model performance in ML [51]. In fact, model performance, according to Schumann et al. [89], can be improved by scaling up datasets, while Birhane et al. found that “scaling up” is one of the top desired values in ML research amongst the top 100 most influential ML papers of the past decade published in two of the most prestigious AI conferences (NeurIPS and ICML)[12].

Scale is furthermore presented as a shortcut that can circumvent various dataset curation related problems such as problematic content, resource-intensive dataset curation, and costly annotation processes, where larger scale is seen as a substitute for quality data and to ensure coverage of long tail of “uncommon” samples. Jia et al., for example, claim that: “*heavy work on data curation and annotation*” can be avoided by scaling up image-text datasets [49]. The “scale beats noise” narrative has tactfully re-framed thoughtful handheld dataset curation as a costly problem that can be “solved” by larger scale. Scale, according to such narrative, is a liberating panacea that not only frees the downstream ML pipeline from the burdens of expensive filtering or post-processing steps but also makes up for “noisy” data as if captioning errors in multimodal datasets of image and alt-text pairs can somehow be averaged out through the correct captioning elsewhere in the dataset. Such lines of thinking are not unique to this specific context, but form a widespread belief that drives initiatives such as the LAION datasets and permeate the entire field of the multi-modal models.

2.2 The Cost of Scale Thinking

The primary motivation behind the LAION-400M undertaking was to produce open-source variants of the opaque Web-Image-Text (WIT) dataset, and the CLIP [78] and DALL.E [80] models. Such open-sourcing initiatives are important first steps towards accountability and building trustworthy AI given that for any auditing and evaluating to take place, open access is a crucial prerequisite. Nonetheless, given the numerous concerns that arise with web-sourced data, continual evaluation and audit of large-scale datasets and models is imperative for well-functioning, just, and healthy open-sourcing practices.

For instance, Science and Technology Studies (STS) scholars and critical data and AI studies have repeatedly emphasized that “scale thinking” stands in stark opposition to values such as societal equity or effective systemic change [39, 53]. In fact, unwavering commitment to scalability is instrumental to the realization of Big Tech’s objectives, such as profit maximization and market monopoly, enabling the centralization of power in the handful few. This often comes at the expense of advancing and cultivating values such as individuals’ rights, informed consent, justice, and consideration for societal impacts of models [12].

With the awareness of the amplified negative downstream impacts of scale, there has also been increased attention towards the need to evaluate and audit models and large-scale datasets as an important intervention and accountability mechanism [79, 65, 99, 60]. The recent emergence of grassroots-based open-sourcing initiatives come as a response to the increasing adoption of the closed-source commercial API access mode of dissemination being used for projects such as GPT-3 [16], CLIP, and DALL.E. For instance, EleutherAI has achieved success by replicating both the *WebText* dataset (on which GPT-3 was trained) and the GPT-3 model itself by carefully curating and disseminating the Pile dataset [30] and training and sharing the GPT-Neo [14]/GPT-NeoX [3] models. In this regard, the open-sourcing movement has been critical, enabling open access to datasets and models, which is key for independent auditing and evaluation.

3 Dataset Audit: LAION-400M and LAION-2B-en

One of the challenges of auditing multimodal datasets is that hateful content, negative stereotypes, and otherwise harmful and marginalizing representations can be present in either modality: text or image. This means that audits can span techniques ranging from image content analysis [92, 13], image source analysis (by analyzing the URL field), image-text cross-modal analysis (looking for discordance between an image and its alt-text description) and alt-text content analysis. Poor data quality, for example, is a common issue that arises with scale. Audits on image content analysis, for example, have revealed that nearly 30% (approximately 700 million image-text pairs) in the LAION-2B-en dataset are duplicates [105]. This, as addressed in [92] and [105], can manifest as *Digital Forgery*, or exact memorization of training examples present multiple times in training data, which was shown to be possible in recent work by Carlini et al [18] – a phenomenon that has stark ramifications for the field of image generation at large.

3.1 Audit Methodology

Our audits are focused on two versions of the LAION visio-linguistic datasets: LAION-400M [88] and LAION-2B-en, the English-language subset of the larger LAION-5B dataset [87] that consists of 2.32 billion image-text pairs. The LAION-400M dataset is currently available as a collection of 32 randomly sampled subsets and are obtainable as individual parquet² files with a mean size of 1.68 GB, for a total of 54 GB. The LAION-2B-en, on the other hand, consists of 128 parquet files, each with a mean size of 2.52 GB (and a total size of 321 GB). Each of these parquet files contains image-text data pertaining to the following data fields: ['SAMPLE_ID', 'URL', 'TEXT', 'HEIGHT', 'WIDTH', 'LICENSE', 'NSFW', 'similarity'].

In order to evaluate the impact of scaling a dataset from 400 million to 2 billion samples on hateful content, we perform the following audits. We first sub-sample the dataset(s) and then extract the alt-text descriptions associated with the sampled image-rows in the 'TEXT' field (See Figure 1). In all our experiments, we randomly sample 0.1 million rows from each of the 160 (= 32 + 128) constituent parquet files spanning the two datasets. This yields $N_{samples,400M} = 3.2$ million samples for the LAION-400M dataset and $N_{samples,2B-en} = 12.8$ million samples for the LAION-2B-en datasets respectively. To this end, we use *Pysentimiento* [72], a SoTA open-source NLP framework.

To begin with, we define the metric, Hate Content Rate (HCR)³: $\psi_{type}(P_{threshold})$ to be,

$$\psi_{type}(P_{threshold}) = 100 \times \frac{\sum_{i=1}^{N_{samples}} \mathbb{1}(\tilde{p}_{type,i} > P_{threshold})}{N_{samples}} \quad (1)$$

where $\mathbb{1}(\cdot)$ is the indicator function, $\tilde{p}_{type,i}$ is the probability score assigned by the Pysentimiento model for the text associated with the i^{th} sample and $type \in \{'hateful', 'targeted', 'aggressive'\}$.

This captures the ratio of samples (as a percentage) that resulted in the Pysentimiento model assigning the associated hate/targeted/aggressive speech probability score to be greater than $P_{threshold}$. We perform a comparative analysis of the extent of *hate speech*, *targeted speech* and *aggressive speech* contained in them using Pysentimiento. In response to an input sentence, *Pysentimiento* outputs a 3×1 vector containing probability scores across the three categories of *hateful*, *targeted* and *aggressive* speech (see Table 1 for randomly selected samples). The extracted alt-text descriptions are then passed through Pysentimiento to extract the $N_{samples} \times 3$ text-quality score matrices for each of the two datasets.

We also introduce the 'Any-of-the-three' detector that maps to the case where the input text fails the quality test if *any* of $\tilde{p}_{hateful}$, $\tilde{p}_{aggressive}$ or $\tilde{p}_{targeted}$ happens to be greater than $P_{threshold}$. The associated 'Any-of-the-three'-

²“Apache Parquet is an open source, column-oriented data file format designed for efficient data storage and retrieval. It provides efficient data compression and encoding schemes with enhanced performance to handle complex data in bulk”

³We use the metric Hate Content Rate (HCR) as a shorthand for not just hateful content but all the three categories: hateful, targeted, and aggressive.




Figure 1: Experimentation details: Dataset sub-sampling, inference using Pysentimiento and thresholding for estimating Hate Content Rate (HCR).

Table 1: Samples of alt text descriptions found in the dataset and the probability scores across the three categories of *hateful*, *targeted* and *aggressive* speech.

Alt text	$P_{hateful}$	$P_{targeted}$	$P_{aggressive}$
'Biden's Spending Will Go To Illegal Immigrants While Tax Hikes Will Destroy American Jobs'	0.902	0.024	0.449
'If you know this man, please, for the love of God tell him to BURN these pants!!'	0.401	0.262	0.517
'shut up and be a don like nancy - Personalised Men's Long Sleeve T-Shirt'	0.395	0.559	0.128
'This bored rich blonde shoplifter gets rough f**keds'	0.934	0.895	0.128
'Horny slave tied to tree gets pulled on her beautiful tits and gets hit on her c*nt with a stick and hands'	0.983	0.911	0.909

HCR, $\bar{\psi}(P_{threshold})$ would be:

$$\begin{aligned}
 \bar{\psi}(P_{threshold}) &= 100 \times \frac{\sum_{i=1}^{N_{samples}} \mathbb{1}((\tilde{p}_{hateful,i} > P_{threshold}) \parallel (\tilde{p}_{targeted,i} > P_{threshold}) \parallel (\tilde{p}_{aggressive,i} > P_{threshold}))}{N_{samples}} \\
 &= 100 \times \frac{\sum_{i=1}^{N_{samples}} \mathbb{1}(\max\{\tilde{p}_{hateful,i}, \tilde{p}_{targeted,i}, \tilde{p}_{aggressive,i}\} > P_{threshold})}{N_{samples}}
 \end{aligned} \tag{2}$$

We then perform a quality check where we evaluate if one (or any) of the 3 probability score values associated with an input alt-text description exceeds a certain pre-set threshold score $P_{threshold}$, in which case the input text is deemed to have failed the quality check at that threshold (Figure 1 illustrates this process). We lastly compare the statistics associated with the text-quality score matrices to understand the nature of the text that was scooped in when the dataset expanded from 400 million samples to 2 billion samples.

4 Dataset Audit Results

In this section, we use both HCR and, more specifically, 'Any-of-the-three'-HCR, $\bar{\psi}(P_{threshold} = 0.5)$ as the default metric of comparison to characterize the amount of problematic content in both LAION-400M and LAION-2B-en datasets. We also and carry out a file-wise comparison of specific shards of both datasets.




Figure 2: HCR curves for the LAION400M and LAION-2B-en datasets using *PySentimiento* outputs. As the dataset is scaled, there is a statistically significant increase in hateful content.

4.1 Scaling is not Benign: Comparing LAION 400M and LAION 2B-en

We begin by focusing on Figure 2 that presents a plot of the HCR curves as a function of $P_{threshold}$. We observe that, as $P_{threshold}$ increases, the HCR curves monotonically decrease, indicating that fewer textual samples meet the more stringent constraint placed by a higher $P_{threshold}$ value. Worryingly however, we found that for all the sentiment types – *hate*, *targeted*, and *aggressive* speech – the HCR-curve(s) pertaining to the 2B-en dataset lies *strictly* above the 400M dataset’s curve(s). This signifies that irrespective of what $P_{threshold}$ is being chosen, the quality failure rate signifying the prevalence of hateful content is higher with the 2B-en dataset in comparison to its 400M counterpart. We found that amongst the three sentiment types, the ‘hateful’ type emerged as the most prevalent for both datasets, with the 2B dataset having a HCR of up to **0.7%** and 400M one of **0.6%**, followed by the ‘targeted’ type, with an HCR up to **0.25%** v/s **0.2%**, and finally the ‘aggressive’ type, with an HCR of **0.04%** v/s **0.03%**.

For the ‘Any-of-the-three’ curve (leftmost in Figure 2), we observe that the HCR curves pertaining to the 2B-en dataset is above the the curves of the 400M dataset. Given that both LAION-400M and LAION-2B-en are extracted from the *CommonCrawl* dataset, we hypothesize that during the race to expand the dataset to 2 billion samples, the dataset scraping module might have sampled from the *low-quality sub-graphs* of the *CommonCrawl* graph at a rate worse than that during the LAION-400M creation process. We also note that the *CLIP-filtering* threshold to have been relaxed from **0.3** (for LAION-400M) to **0.28** (for LAION-2B-en), which could be another explanatory factor.

In order to investigate this phenomenon of increased presence of hateful, targeted and aggressive content with scale deeper, we firstly perform *binomial proportion confidence interval analysis* to establish lower and upper confidence level of ‘Any-of-the-three’-HCR for both datasets at a given reasonable $P_{threshold}$ of **0.5**. For this, we use the *Wilson Score interval method* [107] with coverage at 0.95 (or $\alpha = 0.05$) that resulted in:

$$\begin{aligned} \bar{\psi}(0.5) &\in [\bar{\psi}_{lb,dataset}^{(\alpha=0.05)}(0.5), \bar{\psi}_{ub,dataset}^{(\alpha=0.05)}(0.5)] \\ \bar{\psi}_{400M}(0.5) &= 0.298 \in [0.292, 0.304] \\ \bar{\psi}_{2B-en}(0.5) &= 0.344 \in [0.341, 0.347] \end{aligned} \tag{3}$$




Figure 3: Binomial proportion confidence interval (CI) analysis to establish the extent of HCR underestimation upon using LAION400M statistics.




Figure 4: Fused swarm-box-violinplot that captures the file-wise HCR metrics for all the 160 (=32+128) parquet files from LAION400M and LAION-2B-en.

where $\bar{\psi}_{lb,dataset}^{(\alpha=0.05)}$ and $\bar{\psi}_{ub,dataset}^{(\alpha=0.05)}$ are the lower-bound and the upper-bound values of the confidence interval at $\alpha = 0.05$. As seen in Equation 3, the lower-bound HCR for the 2B-en dataset is markedly **higher** than the upper-bound estimate of HCR for the 400M dataset thus leading to change-of-HCR, $\delta_{CI} = \left(\frac{\bar{\psi}_{lb,2B-en}^{(\alpha=0.05)}(0.5) - \bar{\psi}_{ub,400M}^{(\alpha=0.05)}(0.5)}{\bar{\psi}_{ub,400M}^{(\alpha=0.05)}(0.5)} \right) \times 100$ of 12.26% (See Figure 3). Note that even under this benevolent setting where we compute the difference between the *lower-bound* estimate of HCR for the 2B-en dataset and the *upper-bound* estimate of HCR for the 400M dataset, we still see a 12.26% normalized increase in HCR.

We have so far established the risks of extending the LAION-400M dataset quality statistics to its bigger counterpart, that is LAION-2B-en. This, as we expand further in Section 8, is a consequence of rich non-iid inter-sample correlations emerging from a graph-structured prior for *CommonCrawl*. This begs the question: Given that the dataset was uniformly sampled at the shard/file-level, will the HCR statistics computed at the shard/file-level compare well with the global dataset-level HCR statistic? We investigate this below.

Table 2: The file-wise HCRs for LAION-2B-en are statistically higher than their LAION-400M counterparts. A table capturing results from the two-sample t-test while correcting for unequal variances (using the *Welch separate variances T-test*).

	T	dof	p-val	cohen-d	BF10
hateful	14.48	53.32	2.019874e-20	2.64	3.785e+27
targeted	13.80	54.91	8.443671e-20	2.47	5.957e+25
aggressive	4.44	47.96	2.601226e-05	0.87	2131.144

4.2 Intra-dataset Filewise Comparisons

Given that the two datasets, LAION-400M and LAION-2B-en, are split into 32 and 128 purportedly uniformly sampled shards, respectively, we now examine the validity of the file-level HCR metrics to the global dataset-level metrics. We use the $0.1 \text{ million} \times 3$ sized file-level text-quality score matrices obtained from the *Pysentimiento* model and compute what fraction of these rows are greater than $P_{threshold}$ of 0.5 for all the 3 columns. This yields file-level HCRs (in%) for each of the two datasets with the 3 columns mapping to *hateful* speech, *targeted* speech and *aggressive* speech.

We found that the file-wise HCRs all tightly cluster around the mean levels for the individual datasets. Figure 4 shows the fused swarm-box-violinplot that captures the file-wise HCR metrics for all the 160 (=32+128) parquet files spanning the two datasets. For example, the ‘hateful’ related HCR for LAION-400M has a mean value of **0.298** which **increased to 0.344** for LAION-2B-en. All the 32 constituent file-wise HCRs for this dataset fall within **0.26** and **0.33**. Furthermore, 97% of all the files have their HCRs within 2 standard deviations of the mean-HCR for the dataset. Similarly, the mean-HCR for the entire LAION-2B-en dataset is 0.344 and the range across all the 128 files is (0.297,0.382], that renders 95% of all the constituent files to have their file-level HCRs to be within 2 standard deviations of the dataset level mean HCR.

We also observe that the 128 file-level HCRs for LAION-2B-en (the red swarms) are higher than the 32 file-level HCRs for the LAION400M (the blue swarms) in Figure 4 for all the three sub-categories of hate speech. In order to ascertain if this difference is statistically valid, we perform a two-sample t-test while correcting for unequal variances (using the *Welch separate variances T-test* [26]) and explicitly setting the alternative hypothesis set to be ‘greater’ (with respect to the alternate hypothesis that the mean of the 2B-en HCRs is greater than the mean of 400M-HCRs).

The results of this two-sample t-test are captured in Table 2. As seen, for all the 3 categories of ‘hateful’, ‘targeted’ and ‘aggressive’ speech, the strong T-values **14.48**, **13.8**, and **4.44** combined with high Cohen’s-d (**2.64**, **2.47**, **0.87**) and low p-values (all $\ll 1e^{-4}$) strongly support the hypothesis that the file-wise HCR associated with the 2B-en dataset is **higher** than the file-wise HCR for the 400M dataset, thus adding further evidence to our claim of dataset degradation upon dataset scaling. (Here, ‘dof’: degrees of freedom, ‘BF10’: Bayes Factor of the alternative hypothesis and ‘power’: 1 - type II error:= Achieved power of the test⁴).

5 Model Audit: Scale and Visio-linguistic Bias

In the previous section, we demonstrated that hate content of the image alt-text descriptions increased when the dataset size was increased from 400 million to 2 billion samples. In this section, we examine the downstream consequences of dataset-scaling on CLIP-like visio-linguistic models trained with these dataset variants.

5.1 Audit Methodology

In order to quantitatively evaluate the downstream consequences of problematic dataset on models, we explored model variants where the architecture was held constant and two model checkpoints were being provided: one trained with LAION-400M and the second trained with LAION-2B-en. The emergence of OpenCLIP [46] facilitated this endeavor as (to the best of our knowledge) it remains the only resource that hosts visio-linguistic model variants with *fixed model architecture* but varying dataset sizes (trained on LAION-400M and LAION-2B-en datasets respectively). OpenCLIP (at the time of our experimentation) provided the following CLIP-model pairs presented in Table 3 that met our criteria.

The OpenCLIP project currently uses an idiosyncratic naming convention for the model checkpoints presented in the right column of Table 3 (this is further covered in Appendix A).

⁴See <https://pingouin-stats.org/build/html/generated/pingouin.ttest.html>

Table 3: Architecture-Dataset variants in the OpenCLIP ecosystem that meet our criteria.

Architecture	Dataset/Checkpoint
ViT-B-32	openai
	laion400m_e31
	laion400m_e32
	laion2b_e16
	laion2b_s34b_b79k
ViT-B-16	openai
	laion400m_e31
	laion400m_e32
	laion2b_s34b_b88k
ViT-L-14	openai
	laion400m_e31
	laion400m_e32
	laion2b_s32b_b82k




Figure 5: A sample of images from the Chicago Face Database (CFD) across the 8 **self-identified** race-gender combinations. The titles of each of these images are the exact file-names of these images in the CFD 3.0 version that is hosted at <https://www.chicagofaces.org/download/>.

Amongst all these 13 model-dataset pairs presented in Table 3, we focus on the checkpoints associated with the ViT-L-14 model architecture that the LAION-5B paper [87] presents as the largest (428 million parameters). We present more details about this ViT-L-14 model in Appendix E. Also, given that the ViT-L-14 backbone has two variants ‘laion400m_e31’ and ‘laion400m_e32’ trained on the same LAION-400M dataset (signifying checkpoints derived after 31 and 32 epochs respectively), we chose the *most-trained* checkpoint that is ‘laion400m_e32’. Thus, the three model variants that we experimented with are: [(ViT-L-14, openai), (ViT-L-14, laion400m_e32), (ViT-L-14, laion2b_s32b_b82k)].

In order to evaluate the effect of scaling dataset on these model variants, we used the Chicago Face Dataset (CFD) [61], as a probe dataset. We replicated the *Zero-Shot CLIP experiment* that appeared in *Section 7.1-Bias* of the original CLIP paper [78] by OpenAI, the details of which are in subsection 5.2. The CFD is a highly controlled dataset that consists of high resolution⁵ images of 597 unique individuals along with their *self-identified* race and gender labels belonging to Asian (109), Black (197), Latin (108), and White (183) categories. A sample of images from the CFD dataset is shown in Figure 5. The dataset has been meticulously standardized in order to control for potentially confounding

⁵The images are sized $2444(w) \times 1718(h)$ pixels and “equated for color temperature and placed onto a plain white background”. Of the 597 individuals, 307 self-identified as ‘female’ and ‘290’ self-identified as ‘male’.




Figure 6: Heatmap plots to help the reader visualize the (a) The 8×77 token matrix, (b) The 768×8 text-features matrix, and (c) The 597×768 image-features matrix.

causal covariates such as facial expressions, resolution, image-pixel saturation, lighting conditions, clothing, and eye gaze. The 597 images have each individual wearing the same heather grey t-shirt. While much smaller in volume, unlike the FairFace dataset [52], the individuals in CFD had their consent obtained, were financially compensated and were given the agency to *self-identify* their race and gender categories⁶.

5.2 Experiment Design

The sub-phases involved in the bias analysis experiments (or the *human-being* experiment as we term it hereafter) were as follows:

1: Image pre-processing: All the 597 images with neutral expressions were extracted from CFD were pre-processed using OpenCLIP’s built-in `preprocess` function that entails resizing (to size 224×224), center-cropping and pixel intensity normalization sub-processes. The output of this sub-phase is a CFD-image-tensor, $\mathbf{I}_{cfid} \in \mathbb{R}^{597 \times 224 \times 224 \times 3}$.

2: Class-generation and tokenization: As explained above, we first created an 8-class vector with the classes being [‘human being’, ‘animal’, ‘gorilla’, ‘chimpanzee’, ‘orangutan’, ‘thief’, ‘criminal’ and ‘suspicious person’]. Except the ‘human being’ class, the last 7 classes were verbatim extracted from Section 7.1 Bias of the CLIP paper. Next, we created the class-sentences using “A photo of a/an <class>” template⁷. The output of this sub-phase is a sparse zero-padded text-token matrix, $\mathbf{T}_{8-class} \in \mathcal{I}^{8 \times 77}$ where $\mathcal{I} = [0, \dots, N_{tokens} - 1]$ is the tokenizer-index set (See Figure 6(a) for a heatmap-visualization of this matrix).

3: Forward pass, feature extraction and norming: The pre-processed image tensors and the text-tokens generated in the previous sub-phase were now fed into the encoder of the OpenCLIP model chosen and the output image and text features were then normalized. For the ViT-L-14 model, these are 768-dimensional features thus rendering the text and image feature-matrices over the 597 neutral-expression CFD images to be 597×768 . That is, the image-feature matrix is $\mathbf{F}_I = [\mathbf{f}_0^I, \dots, \mathbf{f}_{596}^I]^T \in \mathbb{R}^{597 \times 768}$ (heatmap in Figure 6(c)) and the text-feature matrix would be: $\mathbf{F}_\tau = [\mathbf{f}_0^t, \dots, \mathbf{f}_7^t]^T \in \mathbb{R}^{768 \times 8}$ (heatmap in Figure 6(b)). In order to highlight how self-similar the 8×8 textual-features are, we present Figure 7(a) that has the annotated heatmap of the $\mathbf{F}_\tau \times \mathbf{F}_\tau^T$ matrix. Similarly, we also present Figure 7(b) that has the heatmap of the 597×597 sized $\mathbf{F}_I \times \mathbf{F}_I^T$ matrix. Given the fact that the 597 images were sorted and grouped by Race-Gender categories, the block-like structures visible in Figure 7(b) indicates the fact that the model’s output image-features are influenced by these categorical indicators.

4: Computing softmax-matrices: Firstly, we obtain the image-text cosine-similarity matrix, $\mathbf{C} \in \mathbb{R}^{597 \times 8}$ by:

$$\mathbf{C} = \mathbf{F}_I \mathbf{F}_\tau^T. \quad (4)$$

Then, the softmax-matrix $\mathbf{S} \in \mathcal{P}^{597 \times 8}$ ($\mathcal{P} = \{p | 0 < p < 1\}$) is computed as:

$$\mathbf{S} = \text{softmax}(100 \times \mathbf{C}). \quad (5)$$

⁶We note that the binary gender category and the seemingly clean race classification is a limitation of the CFD given that genders and races are fluid, complex, multivalent, and multidimensional in actuality. Yet, despite this limitation, we believe the dataset presents a useful proxy in the context of our experiments.

⁷As advocated in the Interacting with CLIP jupyter notebook shared at https://github.com/mlfoundations/open_clip/blob/main/docs/Interacting_with_open_clip.ipynb in the context of Zero-Shot Image Classification for CIFAR-100 dataset. These 8 sentences were then *tokenized* using OpenCLIP’s tokenizer module (the Vocab size is 49408 for all the models considered in this paper) that thus yielded a 8×77 sized token-matrix.




Figure 7: Heatmap plots to help the reader visualize the (a) Sentencified-class self-similarity matrix: $\mathbf{F}_\tau \times \mathbf{F}_\tau^T$ and (b) CFD images self-similarity matrix: $\mathbf{F}_I \times \mathbf{F}_I^T$

Here $\text{softmax}()$ is the softmax function applied row-wise. That is, if $\mathbf{C}_{i,j}$ is the i^{th} row j^{th} column element in the cosine-matrix, then the corresponding $(i,j)^{th}$ element in the softmax-matrix, $\mathbf{S}_{i,j}$ would be $\mathbf{S}_{i,j} = \frac{100 \times \exp(\mathbf{C}_{i,j})}{\sum_{j=0}^7 (100 \times \exp(\mathbf{C}_{i,j}))}$.

In Figure 8, we present the three 597×8 sized output softmax-matrices obtained from the Vit-L/14 family of models, all with the same 428 million parameters and fixed architecture with only the (pre)training dataset being varied across the OpenAI-WIT, LAION-400M and LAION-2B-en choices. The $(i,j)^{th}$ element of each of these matrices captures the output softmax value pertaining to the j^{th} class ($j \in \{0, \dots, 7\}$) obtained from that specific OPEN-CLIP model in response to the i^{th} input CFD image ($i \in \{0, \dots, 596\}$). The 597 rows (representing the 597 CFD images) are grouped by their self-identified Race-Gender groupings. That is, the first 57 rows represent images from the Asian-Female (abbreviated as AF), and the next 52 rows map to the Asian Male (AM) group, and so on. The title of these subplots is formatted as strings with 3 fields separated by the ‘|’ character: [`< cfd_Vit-L-14 > | < training-dataset > | < Phuman >`]. Here, P_{human} is the probability that the top-predicted class (with the highest cosine-similarity/softmax values) is the 0^{th} class mapping to ‘human-being’. That is,

$$P_{human} = \frac{\sum_{i=1}^n \mathbb{1} \left(\text{argmax}_{j \in \{0, \dots, 7\}} \left(\sigma_j^{(i)} \right) = 0 \right)}{n}, \quad (6)$$

where $\sigma_j^{(i)}$ is the softmax score pertaining to the j^{th} class in response to the i^{th} image, $\mathbb{1}(\cdot)$ is the indicator function and $n(= 597)$ is the number of images.

6 Model Audit Results

We found that none of the model variants associated human images from CFD with P_{human} with a high (close to 1) score. Instead, these models yielded a P_{human} score closer to **0.2**. We detail the following observations (see Figure 8).

Observation-1: Both models trained on 400M samples from LAION400M and OpenAI-WIT label images of humans from CFD as one of the racist and dehumanizing classes (as opposed to a ‘human being’), with a **0.186** rate of being labelled as P_{human} with LAION-400M. This further decreased to **0.134** for OpenAI-WIT. In other words, OpenAI-CLIP associates **nearly 87%** of the CFD human-face images with the *7 offensive classes* rather than the human-being class, with a particular stress towards the suspicious person class. Comparing the LAION-400M and OpenAI-WIT models, we find that LAION’s model give images of humans offensive class assignments at a slightly lower rate than the OpenAI-WIT model, thereby not only exposing the limitations of ranking models based on ImageNet-1k-zero-shot accuracy, but also bringing into further focus the contents of the WIT-400 million dataset that still remains




Figure 8: Heatmap plots of the three dataset-dependent 597×8 softmax-matrices obtained from the *human being* experiment.

Table 4: Table summarizing the results of the CFD-Vit-L/14 experiments.

dataset $\downarrow$ \ metric $\rightarrow$	ImageNet-acc	P_{human}	$P_{bm \rightarrow criminal}$	$P_{bf \rightarrow criminal}$
openai	0.753	0.134	0.204	0.221
laion400m_e32	0.739	0.186	0.140	0.212
laion2b_s32b_b82k	0.754	0.094	0.774	0.413

beyond public access.

Observation-2: When the dataset was scaled from 400M samples (LAION-400M) to 2 billion samples (LAION-2B-en), P_{human} fell by nearly half to **0.094**, from **0.186** with most of the softmax-mass being allocated to the *criminal* and *suspicious person* classes.

Observation-3: Another consequence of the dataset scaling from 400M to 2B was the notable shifting of the softmax-mass from the *human being* class to the *criminal* class, especially for the Black-Female (BF) and Black-Male (BM) categories. In order to further ascertain if this was just a visual artifact of the heatmap plot and the association of criminality to faces belonging to the BF/BM categories, we performed a depth-wise analysis.

We found that the mean softmax score for the *criminal* class that the model allocates to Black-female faces *more than doubled* from **0.22** to **0.45** when the dataset was scaled from LAION-400M to LAION-2B-en. Similarly, the mean softmax score for the *criminal* class in CFD *nearly tripled* from **0.22** to **0.65** for Black-male faces with dataset scaling. Figure 9 presents this by means of categorical box-plots of the softmax scores along with the mean and variance statistics in the titles. Furthermore, misclassification rates increased with scale (see Table 3). While **21.2%** of the Black-female faces had a top-predicted class of *criminal* for the LAION-400M model, this number *almost doubled* to **41.3%** for the LAION-2B-en model (This is captured in the $P_{bf \rightarrow criminal}$ column of Table 3). Most notably these misclassification rates for the Black-Male category ($P_{bm \rightarrow criminal}$) *increased nearly by five-fold* from **14%** to **77.4%**.

For further clarity, we summarize these results in Table 4. As we can see, the only metric where we spot the so-termed ‘progress’ is in the ‘ImageNet-acc’ column which maps to the ImageNet-Zero-shot-1k top-1 accuracy metric⁸. To reemphasize the results observed here, we see that scaling the (pre)training dataset from 400 million samples to 2.32 billion samples did result in a *gain* of 1.5% top-1 accuracy on an idiosyncratic task such as ImageNet-1k [24], but it also ended up **halving** the P_{human} , **doubling** $P_{bf \rightarrow criminal}$, **quintupling** $P_{bm \rightarrow criminal}$ classes.

⁸We have reproduced the results verbatim from the LAION-5B [87] and CLIP [78] papers for the ‘ImageNet-acc’ column.




Figure 9: Box plots demonstrating the variation of the softmax values associated with the `criminal` class for the (a) Black women and (b) Black men category images from CFD.

7 Qualitative Analysis: Dehumanization and Criminalization of Black Bodies

In 2015, Google’s Photo app classified photos of Jacky Alciné and his friend (both of whom Black) as “gorillas”. Eight years later in 2023, the problem remains unsolved [36]. The dehumanization of Black bodies through the comparison and classification of Black people with animals, specifically apes, monkeys and orangutans is not a new phenomenon but the historical origins can be traced back to the thirteenth century [48]. European voyagers referred to West Africans as violent savages, uncivilized, beast-like, and even displayed them in zoos. This phenomena of likening people of African descent to non-human primates has been refereed as the “Negro-Ape metaphor” [58, 35].

The characterization of Black people as “animal-like” placed Europeans and North Americans at the top and Africans lower down, in a closer proximity to apes and other primates in such arbitrary (yet deliberately extractive) race hierarchies [48, 84]. The belief that Black people are closer to apes and are less than humans served as justification for numerous historical atrocities including colonialism, slavery, and the Nazi genocide [84, 66]. These dehumanizing depictions of Black people as monkeys, apes and other primates, still remain a common place in contemporary Western societies and can be found in the way soccer players of African descent in Europe are portrayed [35, 96]; the caricatures of the US president Obama as a chimpanzee in magazines such as the New York Post [4]; the racist name calling of Michelle Obama as “Ape in heels” [48]; and the comparisons of U.S. Rep. Maxine Waters to an orangutan [48], to mention but a few examples. Similarly, we find these harmful and dehumanizing depictions of Black people in the large scale multimodal datasets we examined in this paper (see Figure 10 for a sample of images with the “gorilla” label found in the LAION dataset). A rich body of work with STS and critical data and algorithm studies has emphasized the tendency of ML research, tools, and applications to encode and exacerbate societal stereotypes and historical injustice [9, 69, 17, 63]. As presented in Section 6, our findings extend this rich body of work by demonstrating that not only do large ML models encode such historical trend that dehumanizes Black bodies but also, as these models and datasets increase in scale, such dehumanization of Black bodies is further exacerbated.

The depiction of Black people, particularly Black men, saw a gradual shift from “brute” and “docile” to depictions such as “thug”, “criminal”, and “suspicious” [91] as Black men entered the workforce such as in farms and factories. The rise of for-profit prison industrial complex in the 21st century – many prison companies mandating that municipalities have a 90-95% prison occupancy rate – saw an increase targeted association of Black people. and crime [2, 5]. Such stereotypes and racist ideologies have fueled racial violence, criminalization, and mass incarceration of Black men, especially in the US. Black bodies, according to [10], are often perceived as a threat and typecast as “gangster,” “rapist”, and “ghetto”. The “Black-as-criminal” stereotype, subsequently, can result in non-violent acts of Black men being perceived as violent and aggressive while violent acts performed by white men are perceived as unintentional or get attributed to external factors and uncontrollable causes such as mental health [19].




Figure 10: A collage of images from the LAION datasets that had the term gorilla in the alt-text description that were flagged by the Pysentimiento model as hateful. The precise source and the alt-text descriptions are provided in Appendix G. Note: Sub-figures (b) and (f) have been blurred and pixelated by hand by the authors.

Contrary to these racial stereotypes, a robust body of work, especially in the context of the U.S., documents that Black men commit crimes at a far lower rates than whites, while Black people constitute the group that are victims of violent crimes at far higher rates than whites [37, 33]. Innocent Black people, according to [37], are seven-and-a-half times more likely to be convicted of murder than whites and convicted Black people are 80% more likely to be innocent than other convicted murderers. In 2002, Black people were 6 times more likely to be murdered than whites, and this number was much higher during previous decades, where 47% of victims were African Americans during the 1976-2002 period [82]. Conversely, [77] points out, “African-American adults are 5.9 times as likely to be incarcerated than whites” and more likely than whites to be arrested; once arrested, more likely to be convicted; and once convicted, more likely to be incarcerated than whites. Studies on drug use across demographers in the US reveal a similar trend. Although African Americans and whites use illegal drugs at similar rates, Black people are 19 times more likely to be convicted of drug crimes than innocent whites [37, 82]. Erroneous stereotypes have historically (and currently) served to explicitly, implicitly and systematically place Black people, particularly Black men, as “suspects”, “criminal”, or “persons of interests” [91]. Along with past work that has highlighted the risk of models to amplify racial stereotypes [11, 98, 86], our findings confirm this trend. As outlined in Section 6, we observe that current SoTa models encode and exacerbate racial stereotypes. Furthermore, the likelihood of a Black man or a Black woman to be classified as “criminal” and “suspicious person” *increases as the datasets get bigger*. In Figure 11, we demonstrate four examples where two things happen. Firstly, the association of a Black person’s face with ‘A photo of a criminal’ increases with regards to both the cosine-distance and softmax metrics. Secondly, the cosine-distance between the image(s) and the sentence-ified criminal class for each of these examples increases past the 0.28 threshold that’s used as a semanticity qualifier threshold during dataset curation. This implies that if these images had hypothetically surfaced during dataset curation with these offensive criminality-insinuating textual descriptions, the OPENCLIP model filter




Figure 11: Example images of Black individuals from CFD and the tendency of the OPENCLIP models studied to associate them with the “A photo of a criminal” sentence. The first row of the title(s) indicates the file name, the second row indicates the increase in Cosine similarity and the third row indicates the class-wise normalized softmax values.

trained on the 2 billion samples dataset would have accepted them in and the one trained on the 400 million samples dataset would have filtered them out further illustrating the dark side of scale.

8 Discussion and Recommendations

In this paper, we have systematically examined two datasets (LAION 400M and LAION-2B-en) and models trained on them. Contrary to current discourse around scale, our findings reveal that scale exacerbates hateful content and increases the rates of dehumanizing classifications, particularly those of Black women and men. Datasets are not only fundamental to equitable, just, robust, and well performing models but also rigorous evaluation, audit, curation and management of datasets is critical for advancing the field.

We strongly highlight the need to avoid interpreting the empirical results from a reductionist lens where the emphasis is erroneously laid on the specific trivia pertaining to the metrics introduced (such as P_{human} and $P_{bf/bm \rightarrow criminal}$) and model checkpoint variants used. It is evident that the *brittleness* of these models certainly allows for trivially flipping the results to favor another narrative by smartly changing either the choice of labels, the choice of default-class (replacing human being with a synonym for example), the class sentencification template or the model architecture variants (Using Vit-B/16/32 for example). Besides this, we are certain that parameters beyond our control (such as batch-size used during pre-training, choice of tokenizer and number of training epochs used) also played an important role in influencing these results. Instead, what we are conveying through these results is simply this: in spite of making the most templated design choices pertaining to all the aspects of the pipeline, and in spite of verbatim replication of the empirical orchestration straight from the example code notebooks in the official Github repositories, and in spite of using an extremely controlled *easy* probe dataset and class-design, it was verifiably hard to avoid the glaring negative impact on the biases measured that could be directly attributed to dataset scaling.

Below we present a set of observations that we hope the ML community, dataset curators, as well as other stakeholders would find helpful towards advancing not only data curation but also the field as a whole in a manner that is transparent, rigorous, responsible, and accountable.

Compute constraints As models and datasets get ever larger, ML becomes a field that is dominated by (and accessible to) a handful few within tech corporations and elite universities with unconstrained compute resources, crowding out those outside of it. The presence of big tech affiliated influential papers in ML, for example, show increase from 13% in 2008/09 to 47% in 2018/19 [12]. Assembling large scale datasets requires relatively less resources, time and effort compared to auditing, investigating and “cleaning” them. Conversely, big tech corporations and large institutes with abundant compute power assemble these datasets while often the thorough investigation and cleaning up is left to critical scholars with little resources. In this study, we have done as thorough investigation as we can given our relatively limited resources. Through manual investigation, we have come across various issues, such as poor data quality, for instance overwhelming number of images of screenshots. We were unable to perform a thorough analysis to determine a clear estimate of such poor quality data due to the cost of access to image APIs and the huge compute power required to download the datasets in their entirety in order to sift through them. Even when large datasets such as those that we have audited are accessible, getting the compute and tooling necessary for rigorous audit is a challenge. For instance, simply downloading LAION 2B-en requires 6.2TB of storage, with additional compute needed to carry out analyses such as running *pysentimiento*. We encourage corporations and institutes to perform such audits.

However, such self-audits will remain insufficient. Subsequently, we hope – perhaps through a coalition of the larger community, regulatory, and funding bodies – for the cultivation (though incentives) and creation of an ecology that allocates compute resources for independent auditors without access to institutional compute.

Appropriate and consistent metrics Even though many of these datasets are created to train semantic search systems and image generation models that supposedly “democratize art-creation” for the general public (where a great proportion whom are people of diverse gender, ethnicity and race) the metrics used to check if progress is indeed being made by dataset scaling rarely reflects that diversity. While a certain analyses are being made with regards to the risk of biases and ensuing harm in ethics and safety subsections of reports and articles accompanying datasets, the metrics that supposedly measure these harms are never incorporated as part of the model checkpointing process. For instance, in the ALIGN paper [49], the dataset scaling ablation study focused only on two metrics: the MS-COCO zero-shot retrieval accuracy rates (I2T- $R@1$ and T2I- $R@1$) and the ImageNet K-Nearest-neighbor (KNN) $R@1$ rates. In the BASIC model paper ablation Study [74], the authors gauge the impact of increasing the dataset size from 1.7B to 6B by comparing the ImageNet-1k zero-shot top-1 accuracy. Finally, in the LAION-5B paper [87], the authors use the zero-shot top-1 classification accuracy metric, once again on the ImageNet family of datasets (with the distribution shift variants) and a bespoke VTAB+ benchmark spanning 35 classification tasks covering different computer vision datasets. This means that it is difficult for users to meaningfully compare metrics and performance any of these datasets to each other without re-running analyses. Using standardized, meaningful metrics for measuring progress is important to be able to make informed choices when datasets and to ensure that results are comparable and reproducible.

Mind the non-iid assumptions The δ_{CI} of 12.26%, calculated in Section 4.1 above, has important consequences on estimating the number of low-quality samples that either ought to be filtered out or at least re-investigated on account of having failed the text-quality mechanism that we have proposed. This is a direct limitation that emerges from using statistical rubrics built on the iid (Independent and Identically Distributed) samples assumption. The image samples from the *CommonCrawl* are in violation of the iid assumption as the dataset has an underlying graph-structured prior with rich inter-node correlations. In order to further clarify this, we present results from [67], where the host-level *CommonCrawl* web-graph (where both “*hyperlinks and HTTP redirects and link headers are used as edges to span up the graph*”⁹) was revealed to consist of 384 million nodes and 2.47 billion edges with the largest strongly connected component containing 45.2 million (11.7%) nodes. Similarly, the domain graph constructed by aggregating the host graph on the level of pay-level domains (PLDs) using *publicsuffix.org* as ground truth yielded a graph with 90 million nodes and 1.55 billion edges with the largest strongly connected component spanning 36 million or 40% of the nodes. We furthermore note that the body of graph-sampling literature (See [55] and [44]) cautions us about how the summarizing global-metrics obtained by sampling on graphs can be very different from the values obtained with an erroneous i.i.d sampling assumption. This massive margin illustrates that imprudent extrapolating using confidence intervals, especially on datasets with underlying graph structure with rich inter-node correlations such as the *Common Crawl*, where the sample-level iid (Independent and Identically Distributed) assumptions may stand invalidated. We put forward our findings from these audits as a strong reminder of the limitations of the summary auditing statistics obtained using sub-sampling procedures.

Avoid ad-hoc decision making for dataset curation hyper parameters In the *CLIP inference at the post-processing stage* section of the LAION-5B dataset announcement, we encounter the fact that the dataset curators estimated the cosine-similarity between an image and its alt-text description using the ViT B/32 CLIP model and discarded all images with cosine-similarity score of less than the manually set threshold of 0.28. This is a marked departure from the procedure published during the LAION-400M release where the curators stated that “*We use OpenAI’s CLIP model (the ‘ViT-B-32’ version) to compute the image and alt text embeddings. Then we calculate the cosine similarity of both embedding vectors and drop all samples with a similarity below 0.3. We chose this threshold after trying different values and using human evaluations of how well the texts fit the images. Lower values like 0.28 or 0.29 also seemed okay in many cases, but after further inspections, we decided to choose the conservative value of 0.3*”. The reasoning behind this decision is not clear. However, such a decision might have been taken to boost the dataset size past the 5 Billion mark, a pre-mandated milestone perhaps. Given these decision have a significant consequence for dataset quality, we recommend such processes be rigorously justified, well documented, and made transparent a la scientific practices.

Beware of CFD physiognomy Scholars have warned about the rebirth of phrenology and physiognomy via the by-lanes of Computer Vision [94, 93]. Similarly, some of our preliminary investigations that emerged when we dug into the *whyness* of criminality-association of some CFD faces by the models under consideration shows high correlations with metrics such as Facial Width-to-Height Ratio (fWHR) and Cheekbone Prominence that are recorded as metadata

⁹<https://CommonCrawl.org/category/web-graph/>

in the CFD dataset. Well informed and in-depth awareness of this pernicious development as well as mitigation mechanisms against phrenology is crucial. To this end, we encourage future research to build upon this finding by means of a statistical experiment mapping the objective face-measurement-metrics found in ‘Study-1 and Table-1’ of [61] to the model outputs to further investigate the rebirth of phrenology and develop mitigation mechanisms.

***Pysentimiento* idiosyncrasies and limitations** In this paper, we have used the *Pysentimiento* library to perform textual-quality analysis. As an off-the-shelf computational tool, it inherently lacks a nuanced insight on hateful, aggressive, and targeted speech that might be found in, for example, qualitative methods. However, owing to the growing threat of hate speech and toxic speech in online media [21, 28, 62], toxicity classification and hate speech detection have emerged as highly researched topics within NLP (See [20, 32, 47, 75, 103] for systematic surveys). Besides *Pysentimiento*, there exists a wide set of off-the-shelf options including the *roberta-hate-speech-dynabench-r4*¹⁰ [101, 34], *toxic-bert*¹¹-*Detoxify project* [40], *HateSonar* [23] and the *Perspective API*¹² to measure alt-text quality. We hope that the meta-datasets we have generated pertaining to the 16 million samples considered in Section 3 will be used to not just cross-compare the results between these various hate/toxicity-detection approaches. In order to continue rigorous audits and improve multimodal-toxicity detection models, we encourage future work to use these various NLP models to investigate the impact of scale on hateful content but also to recon with and tackle systemic roots of these problems that require rethinking how we approach hate content beyond technical fixes as emphasized in [76].

Dataset sub-sampling: Only for ethics checks? There is an emergent trend within the broad culture of internal audits (self-audits within big corporations and institutes) focusing *subsample-only-for-ethics-auditing* when it comes to handling large datasets, despite the abundant resources at their disposal. As far as training a monetizable model is concerned, scale is deemed a virtue and not a hindrance as exemplified by frequent aggressive crawl-scrape-scoop strategies. On the contrary, scale is deemed as an impediment when it comes to auditing, evaluating, and stress-testing datasets and models for critical concerns including checking for quality of data, encoded racial stereotypes, and bias. For example, we observed that the CLIP model was trained on a black-box Web-Image-Text (WIT) dataset spanning 400 million image text pairs. However, when it came to measuring the racial biases baked into the model, sub-sampling was resorted to a comparatively *small* dataset, the FairFace dataset [52], which only contains 0.027% (108,501 images) of the training dataset size. Moreover, the bias-measurement exercise is minimal, limited only to running inference (read forward pass) through the model that is an order of magnitude less computationally intensive compared to training the model (backward pass). As stated in *Section 7.1: Bias* in the CLIP paper [78], only 10000 images (0.0025% of the training dataset size) were used from this FairFace dataset for the bias-check-inference task (that we have used in our experiments (see Section 5)). We recommend audit, evaluation, and general critical and ethics work is carried out to the highest possible standards and scientific rigour. Otherwise, it risks ethics washing.

Legal and policy implications The multimodal datasets we audited form a crucial backbone for ML systems, including generative AI. These models are not a purely intellectual exercise but are integrated into society directly or indirectly impacting actual people. Subsequently, legal issues arise from multiple angles, including: consent and rights of individuals in datasets, what should be in datasets and how they should be evaluated and maintained, and mechanisms for responsibility and accountability for problematic content in the dataset as well as the downstream effect on models trained on it. Closing access to datasets used for popular and impactful models as well as active obfuscation of information around these datasets present a major obstacle to developing appropriate regulatory guidelines and guardrails. In this audit study, we have presented extensive evidence of exacerbation of hateful content correlated with scale. We hope this work serves as an initial document for legal and policy experts alike that both demystifies multimodal datasets and illustrates the negative implications of scale.

9 Future Work and Conclusion

We have carried out an extensive audit investigating the impact of dataset scale on hate content and the downstream impact of this on visio-linguistic models trained on such datasets. In this regard, the emergence of projects such as *openclip* [46] have been instrumental in allowing us easy orchestration of the type of investigations presented here. This section presents a list of natural extensions of our work.

BLIP and other CLIP models In the associated github repository, we have shared image-class cross-tabulated softmax matrices akin to the ones presented in Figure 8 for the other non-SoTA CLIP models presented in Table 3

¹⁰<https://huggingface.co/spaces/evaluate-measurement/toxicity>

¹¹<https://github.com/unitaryai/detoxify>

¹²https://developers.perspectiveapi.com/s/?language=en_US

for which we could run the *fix-architecture-vary-training-datasets* experiments presented in Section 5. We highly encourage for these experiments to be replicated across the other models including BLIP [56] and the new variants emerging on the scene. We hope that this will help the ML community to intimately understand (and mitigate) the role that model architectures play in encoding harmful biases as the dataset scales.

Choice of prompt template and class design In this paper, we converted the categorical class labels into sentences using the format “A *photo of* <class>” in order to maintain consistency with the CLIP [78] paper results. We posit that varying this prompt template with its rephrased variants such as “This a *picture of* <class>” would result in variations of the results shown in Section 5. Similarly, we also expect that replacing the word *person* with the self-declared race-gender identifier (such as *asian-man*) will also result in variations to the cosine-similarity value output by the models under consideration. Accordingly, future research might unearth the *fairness-optimal* prompt template by both paraphrasing as well as choosing alternative-identifiers for the word *human being*.

Extension across other expressions and other face datasets In this paper, we have restricted our experimentation to the neutral expression images of the CFD dataset for the sake of brevity. One future avenue for future work might be to investigate if holding the individuals’ faces constant and varying the facial expressions makes a marked difference in the results. Also, inspired by the CFD project, we have seen the emergence of other similar datasets such as MR2 [95], Bogazici face database [85], the Delaware dataset [64] and the ISIEA dataset [109]. Replicating these experiments using these datasets might yield a more granular view of how these models – supposedly trained on internet sourced data – function and what biases might be baked into them.

The Race-Gender experiment: Some initial results There also emerges the natural question with regards to the extent to which stereotypes about facial appearances are cross-related with racial identities by these visiolinguistic models. Given that the CFD has self-identified race-gender labels, we also performed a small scale race-gender classification experiment (similar to the FairFace experiment in the CLIP paper [78]), using the subjects’ self-identified race-gender labels. That is, we replaced the 8 classes of [human being,...,suspicious person] in the *human-being experiment* above with the 8 self-identified race-gender category labels [asian man,...,white woman]. The initial results are discussed in Appendix F and it appears as if faces with visible epicanthic folds (that occurs across a broad spectrum of racial identities) are solely associated with the ‘Asian’ race identifier. This observation merits a deeper analysis especially given the wide availability of meta-data that is associated with the images in CFD that can be a rich source of confounding factors.

9.1 Conclusion

We have carried out a dataset audit of two visio-linguistic multimodal datasets, LAION-400M and LAION 2B-en, and models trained on them. We found evidence of hateful, aggressive, and targeted content in the alt text audit and evidence of racist stereotyping and dehumanizing classification in the models, particularly towards Black men, all of which exacerbates with dataset size. We cannot stress the importance of open-source in audit endeavors such as ours, since any kind of quantitative and qualitative dataset exploration hinges upon access to the artifacts themselves. We are saddened to see an increasing number of ML organizations fail to provide access to their datasets and models, since we believe that this is an essential element to scientific advancement and a healthy, equitable, and innovative research community.

Today’s state-of-the-art visio-linguistic multimodal models are trained with massive carbon footprints, massive data-infrastructure and massive funding. These models are currently being deployed in real-world including in recommendation systems, information-retrieval systems, semantic search systems and image captioning systems, although as we have illustrated in this paper, they can fail at associating photographs of human faces with the description: “A *photo of a human being*”. Given that such failures can result in dire consequences on real people, often those at the margins of society, we implore the research community as well as those developing and deploying these systems to carry out due diligence with rigorous audits of these models and training datasets and take necessary actions, including refraining from use in high-stake scenarios.

Acknowledgements

We would like to thank André Brock, Ellen Rushe, Gary Marcus, Sasha Luccioni, and Thomas Laurent for the invaluable comments on an earlier version of the paper.

References

- [1] ABID, A., FAROOQI, M., AND ZOU, J. Persistent anti-muslim bias in large language models. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society* (2021), pp. 298–306.
- [2] ALEXANDER, M. *The new Jim Crow: Mass incarceration in the age of colorblindness*. The New Press, 2020.
- [3] ANDONIAN, A., BIDERMAN, S., BLACK, S., GALI, P., GAO, L., HALLAHAN, E., LEVY-KRAMER, J., LEAHY, C., NESTLER, L., PARKER, K., PIELER, M., PUROHIT, S., SONGZ, T., WANG, P., AND WEINBACH, S. GPT-NeoX: Large scale autoregressive language modeling in pytorch, 2021.
- [4] APEL, D. Just joking? chimps, obama and racial stereotype. *Journal of Visual Culture* 8, 2 (2009), 134–142.
- [5] BARDES, J. K. Redefining vagrancy: Policing freedom and disorder in reconstruction new orleans, 1862–1868. *Journal of Southern History* 84, 1 (2018), 69–112.
- [6] BASU, A., BABU, R. V., AND PRUTHI, D. Inspecting the geographical representativeness of images from text-to-image models, 2023.
- [7] BENDER, E. M., GEBRU, T., MCMILLAN-MAJOR, A., AND SHMITCHELL, S. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (2021), pp. 610–623.
- [8] BENDER, E. M., AND KOLLER, A. Climbing towards nlu: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics* (2020), pp. 5185–5198.
- [9] BENJAMIN, R. *Race after technology: Abolitionist tools for the new jim code*. John Wiley & Sons, 2019.
- [10] BEY, M. “bring out your dead” understanding the historical persistence of the criminalization of black bodies. *Cultural Studies? Critical Methodologies* 16, 3 (2016), 271–277.
- [11] BIANCHI, F., KALLURI, P., DURMUS, E., LADHAK, F., CHENG, M., NOZZA, D., HASHIMOTO, T., JURAFSKY, D., ZOU, J., AND CALISKAN, A. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale. *arXiv preprint arXiv:2211.03759* (2022).
- [12] BIRHANE, A., KALLURI, P., CARD, D., AGNEW, W., DOTAN, R., AND BAO, M. The values encoded in machine learning research. *arXiv preprint arXiv:2106.15590* (2021).
- [13] BIRHANE, A., PRABHU, V. U., AND KAHSEMBWE, E. Multimodal datasets: misogyny, pornography, and malignant stereotypes. *arXiv preprint arXiv:2110.01963* (2021).
- [14] BLACK, S., LEO, G., WANG, P., LEAHY, C., AND BIDERMAN, S. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow, Mar. 2021. If you use this software, please cite it using these metadata.
- [15] BLODGETT, S. L., BAROCAS, S., DAUMÉ III, H., AND WALLACH, H. Language (technology) is power: A critical survey of “bias” in nlp. *arXiv preprint arXiv:2005.14050* (2020).
- [16] BROWN, T. B., MANN, B., RYDER, N., SUBBIAH, M., KAPLAN, J., DHARIWAL, P., NEELAKANTAN, A., SHYAM, P., SASTRY, G., ASKELL, A., ET AL. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165* (2020).
- [17] BROWNE, S. *Dark matters: On the surveillance of blackness*. Duke University Press, 2015.
- [18] CARLINI, N., HAYES, J., NASR, M., JAGIELSKI, M., SEHWAG, V., TRAMER, F., BALLE, B., IPPOLITO, D., AND WALLACE, E. Extracting training data from diffusion models. *arXiv preprint arXiv:2301.13188* (2023).
- [19] CHAPPLE, R. L., JACINTO, G. A., HARRIS-JACKSON, T. N., AND VANCE, M. Do# blacklivesmatter? implicit bias, institutional racism and fear of the black body. *Ralph Bunche Journal of Public Affairs* 6, 1 (2017), 2.
- [20] CHHABRA, A., AND VISHWAKARMA, D. K. A literature survey on multimodal and multilingual automatic hate speech identification. *Multimedia Systems* (2023), 1–28.
- [21] CITRON, D. K. *Hate crimes in cyberspace*. Harvard University Press, 2014.
- [22] DAVIDSON, T., BHATTACHARYA, D., AND WEBER, I. Racial bias in hate speech and abusive language detection datasets. *arXiv preprint arXiv:1905.12516* (2019).
- [23] DAVIDSON, T., WARMSLEY, D., MACY, M., AND WEBER, I. Automated hate speech detection and the problem of offensive language. In *Proceedings of the international AAAI conference on web and social media* (2017), vol. 11, pp. 512–515.

- [24] DENG, J., DONG, W., SOCHER, R., LI, L.-J., LI, K., AND FEI-FEI, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* (2009), pp. 248–255.
- [25] EDWARDS, B. Flooded with ai-generated images, some art communities ban them completely — ars technica. <https://arstechnica.com/information-technology/2022/09/flooded-with-ai-generated-images-some-art-communities-ban-them-completely/>, Sep 2022. (Accessed on 04/22/2023).
- [26] EFRON, B., AND HASTIE, T. *Computer age statistical inference, student edition: algorithms, evidence, and data science*, vol. 6. Cambridge University Press, 2021.
- [27] ERIC SHERIDAN, K. R. Are we on the cusp of a generative ai revolution? <https://www.goldmansachs.com/insights/podcasts/episodes/02-21-2023-sheridan-rangan.html>, Feb 2023. (Accessed on 04/21/2023).
- [28] FOXMAN, A. H., AND WOLF, C. *Viral hate: Containing its spread on the Internet*. Macmillan, 2013.
- [29] FRY, C. L., NAUGLE JR, T. C., COLE, S. A., GELFOND, J., CHITTOOR, G., MARIANI, A. F., GOROS, M. W., HAIK, B. G., AND VORUGANTI, V. S. The latino eyelid: anthropometric analysis of a spectrum of findings. *Ophthalmic plastic and reconstructive surgery* 33, 6 (2017), 440.
- [30] GAO, L., BIDERMAN, S., BLACK, S., GOLDING, L., HOPPE, T., FOSTER, C., PHANG, J., HE, H., THITE, A., NABESHIMA, N., ET AL. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027* (2020).
- [31] GARCIA, N., HIROTA, Y., WU, Y., AND NAKASHIMA, Y. Uncurated image-text datasets: Shedding light on demographic bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), pp. 6957–6966.
- [32] GARG, T., MASUD, S., SURESH, T., AND CHAKRABORTY, T. Handling bias in toxic speech detection: A survey. *ACM Computing Surveys* (2022).
- [33] GASTON, S. Enforcing race: A neighborhood-level explanation of black–white differences in drug arrests. *Crime & Delinquency* 65, 4 (2019), 499–526.
- [34] GEHMAN, S., GURURANGAN, S., SAP, M., CHOI, Y., AND SMITH, N. A. Realtotoxicityprompts: Evaluating neural toxic degeneration in language models. *arXiv preprint arXiv:2009.11462* (2020).
- [35] GOFF, P. A., JACKSON, M. C., DI LEONE, B. A. L., CULOTTA, C. M., AND DITOMASSO, N. A. The essence of innocence: consequences of dehumanizing black children. *Journal of personality and social psychology* 106, 4 (2014), 526.
- [36] GRANT, N., AND HILL, K. Google’s photo app still can’t find gorillas. and neither can apple’s., 2023.
- [37] GROSS, S. R., POSSLEY, M., OTTERBOURG, K., STEPHENS, K., PAREDES, J., AND O’BRIEN, B. Race and wrongful convictions in the united states 2022. *Available at SSRN 4245863* (2022).
- [38] GROSS, T. F. Own-ethnicity bias in the recognition of black, east asian, hispanic, and white faces. *Basic and Applied Social Psychology* 31, 2 (2009), 128–135.
- [39] HANNA, A., AND PARK, T. M. Against scale: Provocations and resistances to scale thinking. *arXiv preprint arXiv:2010.08850* (2020).
- [40] HANU, L., AND UNITARY TEAM. Detoxify. Github. <https://github.com/unitaryai/detoxify>, 2020.
- [41] HEAVEN, W. D. Generative ai is changing everything, but what’s left when the hype is gone? <https://www.technologyreview.com/2022/12/16/1065005/generative-ai-revolution-art/>, Dec 2022. (Accessed on 04/21/2023).
- [42] HESTNESS, J., NARANG, S., ARDALANI, N., DIAMOS, G., JUN, H., KIANINEJAD, H., PATWARY, M., ALI, M., YANG, Y., AND ZHOU, Y. Deep learning scaling is predictable, empirically. *arXiv preprint arXiv:1712.00409* (2017).
- [43] HOFFMANN, J., BORGEAUD, S., MENSCH, A., BUCHATSKAYA, E., CAI, T., RUTHERFORD, E., CASAS, D. D. L., HENDRICKS, L. A., WELBL, J., CLARK, A., ET AL. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556* (2022).
- [44] HU, P., AND LAU, W. C. A survey and taxonomy of graph sampling. *arXiv preprint arXiv:1308.5865* (2013).
- [45] HUANG, H. The generative ai revolution has begun—how did we get here? — ars technica. <https://arstechnica.com/gadgets/2023/01/the-generative-ai-revolution-has-begun-how-did-we-get-here/>, Jan 2023. (Accessed on 04/21/2023).

- [46] ILHARCO, G., WORTSMAN, M., WIGHTMAN, R., GORDON, C., CARLINI, N., TAORI, R., DAVE, A., SHANKAR, V., NAMKOONG, H., MILLER, J., HAJISHIRZI, H., FARHADI, A., AND SCHMIDT, L. Openclip, July 2021. If you use this software, please cite it as below.
- [47] JAHAN, M. S., AND OUSSALAH, M. A systematic review of hate speech automatic detection using natural language processing. *arXiv preprint arXiv:2106.00742* (2021).
- [48] JARDINA, A., AND PISTON, S. Hiding in plain sight: Dehumanization as a foundation of white racial prejudice. *Sociology Compass* 15, 9 (2021), e12913.
- [49] JIA, C., YANG, Y., XIA, Y., CHEN, Y.-T., PAREKH, Z., PHAM, H., LE, Q. V., SUNG, Y., LI, Z., AND DUEBIG, T. Scaling up visual and vision-language representation learning with noisy text supervision. *arXiv preprint arXiv:2102.05918* (2021).
- [50] JOSHUA LU, R. G. The generative ai revolution will enable anyone to create games — andreessen horowitz. <https://a16z.com/2023/03/17/the-generative-ai-revolution/>, Mar 2023. (Accessed on 04/21/2023).
- [51] KAPLAN, J., MCCANDLISH, S., HENIGHAN, T., BROWN, T. B., CHESSE, B., CHILD, R., GRAY, S., RADFORD, A., WU, J., AND AMODEI, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [52] KÄRKKÄINEN, K., AND JOO, J. Fairface: Face attribute dataset for balanced race, gender, and age. *arXiv preprint arXiv:1908.04913* (2019).
- [53] KOCH, B., DENTON, E., HANNA, A., AND FOSTER, J. G. Reduced, reused and recycled: The life of a dataset in machine learning research. *arXiv preprint arXiv:2112.01716* (2021).
- [54] LEE, T. B. Stable diffusion copyright lawsuits could be a legal earthquake for ai — ars technica. <https://arstechnica.com/tech-policy/2023/04/stable-diffusion-copyright-lawsuits-could-be-a-legal-earthquake-for-ai/>, Apr 2023. (Accessed on 04/22/2023).
- [55] LESKOVEC, J., AND FALOUTSOS, C. Sampling from large graphs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining* (2006), pp. 631–636.
- [56] LI, J., LI, D., XIONG, C., AND HOI, S. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. *arXiv preprint arXiv:2201.12086* (2022).
- [57] LOMAS, N. Shutterstock to integrate OpenAI’s DALL-E 2 and launch fund for contributor artists, 2022.
- [58] LOTT, T. Racist discourse and the negro-ape metaphor. *The Invention of Race: Black Culture and the Politics of Representation* (1999), 7–13.
- [59] LUCCIONI, A. S., AKIKI, C., MITCHELL, M., AND JERNITE, Y. Stable bias: Analyzing societal representations in diffusion models. *arXiv preprint arXiv:2303.11408* (2023).
- [60] LUCCIONI, A. S., AND VIVIANO, J. D. What’s in the box? an analysis of undesirable content in the common crawl corpus. *arXiv preprint arXiv:2105.02732* (2021).
- [61] MA, D. S., CORRELL, J., AND WITTENBRINK, B. The chicago face database: A free stimulus set of faces and norming data. *Behavior research methods* 47 (2015), 1122–1135.
- [62] MANTILLA, K. *Gendertrolling: How misogyny went viral: How misogyny went viral*. ABC-CLIO, 2015.
- [63] MCQUILLAN, D. *Resisting AI: an anti-fascist approach to artificial intelligence*. Policy Press, 2022.
- [64] MENDE-SIEDLECKI, P., QU-LEE, J., LIN, J., DRAIN, A., AND GOHARZAD, A. The delaware pain database: A set of painful expressions and corresponding norming data. *Pain reports* 5, 6 (2020).
- [65] METAXA, D., PARK, J. S., ROBERTSON, R. E., KARAHALIOS, K., WILSON, C., HANCOCK, J., SANDVIG, C., ET AL. Auditing algorithms: Understanding algorithmic systems from the outside in. *Foundations and Trends® in Human–Computer Interaction* 14, 4 (2021), 272–344.
- [66] MONTAGU, M. A. The genetical theory of race, and anthropological method. *American Anthropologist* (1942), 369–375.
- [67] NAGEL, S. Common crawl’s first in-house web graph – common crawl. <https://commoncrawl.org/2017/05/hostgraph-2017-feb-mar-apr-crawls/>, May 2017. (Accessed on 03/15/2023).
- [68] NAKAMURA, L., DAVÉ, S., NISHIME, L., AND OREN, T. G. ‘alllooksamer’? mediating asian american visual cultures of race on the web. *East main street: Asian American popular culture* (2005), 262–272.
- [69] NOBLE, S. U. Algorithms of oppression. In *Algorithms of oppression*. New York University Press, 2018.

- [70] ORSELLI, B. Stable diffusion ai has mastered the female form - niche gamer. <https://nichegamer.com/stable-diffusion-ai-has-mastered-the-female-form/>, Oct 2022. (Accessed on 04/22/2023).
- [71] PACHECO JR, G. *Rhetoric with humor: An analysis of Hispanic/Latino comedians' uses of humor*. The University of Southern Mississippi, 2008.
- [72] PÉREZ, J. M., GIUDICI, J. C., AND LUQUE, F. pysentimiento: A python toolkit for sentiment analysis and socialnlp tasks. *arXiv preprint arXiv:2106.09462* (2021).
- [73] PFOTENHAUER, S., LAURENT, B., PAPAGEORGIOU, K., AND STILGOE, J. The politics of scaling. *Social Studies of Science* 52, 1 (2022), 3–34.
- [74] PHAM, H., DAI, Z., GHIASI, G., KAWAGUCHI, K., LIU, H., YU, A. W., YU, J., CHEN, Y.-T., LUONG, M.-T., WU, Y., ET AL. Combined scaling for open-vocabulary image classification. *arXiv e-prints* (2021), arXiv-2111.
- [75] POLETO, F., BASILE, V., SANGUINETTI, M., BOSCO, C., AND PATTI, V. Resources and benchmark corpora for hate speech detection: a systematic review. *Language Resources and Evaluation* 55 (2021), 477–523.
- [76] PRABHAKARAN, V., WASEEM, Z., AKIWOWO, S., AND VIDGEN, B. Online abuse and human rights: Woah satellite session at rightscon 2020. In *Proceedings of the Fourth Workshop on Online Abuse and Harms* (2020), pp. 1–6.
- [77] PROJECT, S. Report to the united nations on racial disparities in the us criminal justice system.
- [78] RADFORD, A., KIM, J. W., HALLACY, C., RAMESH, A., GOH, G., AGARWAL, S., SASTRY, G., ASKELL, A., MISHKIN, P., CLARK, J., ET AL. Learning transferable visual models from natural language supervision. *arXiv preprint arXiv:2103.00020* (2021).
- [79] RAJI, I. D., AND BUOLAMWINI, J. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial ai products. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society* (2019), pp. 429–435.
- [80] RAMESH, A., PAVLOV, M., GOH, G., GRAY, S., VOSS, C., RADFORD, A., CHEN, M., AND SUTSKEVER, I. Zero-shot text-to-image generation. *arXiv preprint arXiv:2102.12092* (2021).
- [81] ROMBACH, R., BLATTMANN, A., LORENZ, D., ESSER, P., AND OMMER, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 10684–10695.
- [82] ROSICH, K. J. Race, ethnicity, and the criminal justice system.
- [83] SAHARIA, C., CHAN, W., SAXENA, S., LI, L., WHANG, J., DENTON, E. L., GHASEMPOUR, K., GONTIJO LOPES, R., KARAGOL AYAN, B., SALIMANS, T., ET AL. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems* 35 (2022), 36479–36494.
- [84] SAINI, A. *Superior: the return of race science*. Beacon Press, 2019.
- [85] SARIBAY, S. A., BITEN, A. F., MERAL, E. O., ALDAN, P., TRÉBICKÝ, V., AND KLEISNER, K. The bogazici face database: Standardized photographs of turkish faces with supporting materials. *PloS one* 13, 2 (2018), e0192018.
- [86] SCHEUERMAN, M. K., WADE, K., LUSTIG, C., AND BRUBAKER, J. R. How we’ve taught algorithms to see identity: Constructing race and gender in image databases for facial analysis. *Proceedings of the ACM on Human-computer Interaction* 4, CSCW1 (2020), 1–35.
- [87] SCHUHMANN, C., BEAUMONT, R., VENCU, R., GORDON, C., WIGHTMAN, R., CHERTI, M., COOMBES, T., KATTA, A., MULLIS, C., WORTSMAN, M., ET AL. Laion-5b: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402* (2022).
- [88] SCHUHMANN, C., VENCU, R., BEAUMONT, R., KACZMARCZYK, R., MULLIS, C., KATTA, A., COOMBES, T., JITSEV, J., AND KOMATSUZAKI, A. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114* (2021).
- [89] SCHUHMANN, C., VENCU, R., BEAUMONT, R., KACZMARCZYK, R., MULLIS, C., KATTA, A., COOMBES, T., JITSEV, J., AND KOMATSUZAKI, A. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114* (2021).
- [90] SEELOS, C., AND MAIR, J. *Innovation and scaling for impact: How effective social enterprises do it*. Stanford university press, 2017.

- [91] SMILEY, C., AND FAKUNLE, D. From “brute” to “thug:” the demonization and criminalization of unarmed black male victims in america. *Journal of human behavior in the social environment* 26, 3-4 (2016), 350–366.
- [92] SOMEPELLI, G., SINGLA, V., GOLDBLUM, M., GEIPING, J., AND GOLDSTEIN, T. Diffusion art or digital forgery? investigating data replication in diffusion models. *arXiv preprint arXiv:2212.03860* (2022).
- [93] SPANTON, R. W., AND GUEST, O. Measuring trustworthiness or automating physiognomy? a comment on safra, chevallier, gr\ezes, and baumard (2020). *arXiv preprint arXiv:2202.08674* (2022).
- [94] STARK, L., AND HUTSON, J. Physiognomic artificial intelligence. *Fordham Intell. Prop. Media & Ent. LJ* 32 (2021), 922.
- [95] STROHMINGER, N., GRAY, K., CHITUC, V., HEFFNER, J., SCHEIN, C., AND HEAGINS, T. B. The mr2: A multi-racial, mega-resolution database of facial stimuli. *Behavior research methods* 48 (2016), 1197–1204.
- [96] THOMPSON, W. When the beautiful game turns ugly. *ESPN The Magazine* 6 (2013).
- [97] TOMASEV, N., MAYNARD, J. L., AND GABRIEL, I. Manifestations of xenophobia in ai systems, 2022.
- [98] VAN MILTENBURG, E. Stereotyping and bias in the flickr30k dataset. *arXiv preprint arXiv:1605.06083* (2016).
- [99] VECCHIONE, B., LEVY, K., AND BAROCAS, S. Algorithmic auditing and social justice: Lessons from the history of audit studies. In *Equity and Access in Algorithms, Mechanisms, and Optimization*. 2021, pp. 1–9.
- [100] VIDGEN, B., THRUSH, T., WASEEM, Z., AND KIELA, D. Learning from the worst: Dynamically generated datasets to improve online hate detection. *arXiv preprint arXiv:2012.15761* (2020).
- [101] VIDGEN, B., THRUSH, T., WASEEM, Z., AND KIELA, D. Learning from the worst: Dynamically generated datasets to improve online hate detection. In *ACL* (2021).
- [102] VILLALOBOS, P. Scaling laws literature review, 2023. Accessed: 2023-4-22.
- [103] WANG, K., YANG, J., AND WU, H. A survey of toxic comment classification methods. *arXiv preprint arXiv:2112.06412* (2021).
- [104] WASEEM, Z., DAVIDSON, T., WARMSLEY, D., AND WEBER, I. Understanding abuse: A typology of abusive language detection subtasks. *arXiv preprint arXiv:1705.09899* (2017).
- [105] WEBSTER, R., RABIN, J., SIMON, L., AND JURIE, F. On the de-duplication of laion-2b. *arXiv preprint arXiv:2303.12733* (2023).
- [106] WEIDINGER, L., MELLOR, J., RAUH, M., GRIFFIN, C., UESATO, J., HUANG, P.-S., CHENG, M., GLAESE, M., BALLE, B., KASIRZADEH, A., ET AL. Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359* (2021).
- [107] WILSON, E. B. Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association* 22, 158 (1927), 209–212.
- [108] YU, J., XU, Y., KOH, J. Y., LUONG, T., BAID, G., WANG, Z., VASUDEVAN, V., KU, A., YANG, Y., AYAN, B. K., ET AL. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789* (2022).
- [109] ZHENG, Z., LI, S., MO, L., CHEN, W., AND ZHANG, D. Isiea: An image database of social inclusion and exclusion in young asian adults. *Behavior Research Methods* (2021), 1–13.

A OpenCLIP checkpoint naming conventions

In Table 3, openai refers to OpenAI’s closed WIT-400 million samples dataset, laion400m_e31(32) refers to the checkpoint of the model trained on LAION-400M check-pointed after 31(32) training epochs, s32b(s34b) refers to checkpoints where the training was stopped after the model had *seen* 32 (34) billion samples, b79k(82k) (88k) refers to training batch-sizes used (that is batch_size=79000, 82000 and 88000 respectively). This information was not found in the documentation and was gleaned via a github issue raised. The conversation can be found here: https://github.com/mlfoundations/open_clip/issues/454#issuecomment-1451321921. We also note that in the absence of a standardized model-naming template, it is hard to decipher details such as the batch size used for training certain models (especially those named laion400m_e31/e32), which could potentially be another confounding parameter influencing the results obtained in this paper.

B The origins of the dataset scaling laws: A cartoon sketch emerges

While attempting to unearth what this specific dataset scaling law was that the practitioners were so inspired by, we repeatedly encountered a certain cartoon sketch ‘power-law’ plot referred to in both personal exchanges as well as in surveys such as [102]. As it turns out, this cartoon sketch power-plot first appeared as Figure 6 in “*Deep learning scaling is predictable, empirically*” [42], a work that emerged out of Baidu research in 2017. The authors that first presented this plot posit that the generalization error associated with a ML model exhibits a three phase behavior with regards to its training dataset size. The first phase, they state maps to the ‘small data region’, where “*models will struggle to learn from a small number of training samples*” resulting in high generalization errors. The second phase (or the middle portion of learning curves), they claim is the ‘power-law region’, where the generalization error monotonically decreases with training dataset size (linear with application-specific slopes when plotted on a log-log scale). This phase stretches till we hit of point of the ‘glass-ceiling’ or ‘unbreachable error-floor’ on account of factors such as model mismatch and mislabeled data (constituting the third phase). This, of course, has been further supplanted by the likes of the Chinchilla scaling laws (20 tokens per model parameter) [43] in the specialized context of LLMs.

C Blackbox non-reproducible empirical results

As for the blackbox non-reproducible empirical results that validated the dataset-scaling mandate and championed the *scale-beats-noise* narrative, we refer to the ALIGN paper [49] that emerged in 2021. In the abstract section of this paper, we first encounter the following claim: “*We show that the scale of our corpus can make up for its noise and leads to state-of-the-art representations even with such a simple learning scheme*”. The demonstration of this claim appears later in "Section 6.2. Pre-training Datasets" where the authors state that “*To understand better how data size scaling wins over the increased noise, we further randomly sample 3M, 6M, and 12M ALIGN training data and compare them with the cleaned CC-3M data on B7+BERT-base model. Table 10 (sic) shows that while the ALIGN data performs much worse than CC data with the same size (3M), the model quality trained on 6M and 12M ALIGN data rapidly catches up. Despite being noisy, ALIGN data outperforms Conceptual Captions **with only 4x size.***” We note that these experiments (or similar ones) have not been replicated elsewhere to check if these scaling-ratios presented *ipse dixit* in these contexts indeed hold true at all.

D The tactical template: Fuzzy main section meets non-existent appendices

What unites the marquee projects of Dall-E, Parti and Imagen is the near-same tactical template deployed when it comes to (non)declaring the training dataset information. The template runs something like this:

Step-1: Allocate a small nondescript subsection of the main section of the paper covering only the bare minimum details about the number of samples in the training dataset with cross-references to other similar blackbox datasets such as JFT. This coincidentally happens to be Section 4.1 in both Parti and Imagen papers (See Figure 12).

Step-2: Declare that somewhere in the succeeding sections titled on the lines of broader impacts or societal impacts, are details about the ‘potentially problematic’ aspects of the dataset and the downstream risks while patronisingly citing previously published audit papers (such as [13] that have actually done the grunt work of exposing the gory details of such datasets. This happens to be Section 8 - Broader impacts in Parti and Section 6 for the Imagen model.

Step-3: Setting the reader up for a non-existent Appendix section that is not part of the main-paper and not containing any details about how the dataset is actually constructed and where the data is sourced from while noting the fact that its not mandatory for the reviewers to even glance at the Appendix section in peer-reviewed avenues of publishing.

Scaling Autoregressive Models for Content-Rich Text-to-Image Generation

4.1 Training Datasets

We train on a combination of image-text datasets for all Parti models. The data includes the publicly available LAION-400M dataset [43]; FIT400M, a filtered subset of the full 1.8 billion examples used to train the ALIGN model [9]; JFT-4B dataset [44], which has images with text annotation labels. For textual descriptions of JFT, we randomly switch between the original labels as text (concatenated if an image has multiple labels) or machine-generated captions from a SimVLM model [45]. We discuss the limitations of the data in Section 8. For all image inputs, we follow the DALL-E dVAE input processing (Section A.2. Training in [2]) for image tokenizer training and the DALL-E Transformer input processing (Section B.2. Training in [2]) for encoder-decoder training.

8 Broader Impacts

Bias and safety. Text-to-image generation models like GLIDE, DALL-E 2, Imagen, Make-a-Scene, CogView and Parti are all trained on large, often noisy, image-text datasets that are known to contain biases regarding people of different backgrounds. This is particularly highlighted in Birhane et al's [100] analysis of the LAION-400M dataset [43]: their study of the dataset surfaced many problems with respect to stereotyping, pornography, violence and more. Other biases include stereotypical representations of people described as lawyers, flight attendants, homemakers, and so on. Models trained on such data without mitigation strategies thus risk reflecting and scaling up the underlying problems. Our primary training data is selected and highly filtered to minimize the presence of NSFW content; however, we incorporated LAION-400M during finetuning with classifier-free guidance – this improved model performance but also led to generation of NSFW images in some contexts. Other biases include those introduced by the use of examples that primarily have English texts and may be biased to certain areas of the world. In informal testing, we have noticed, for example, that prompts mentioning wedding clothes seem to produce images biased towards stereotypically female and Western attire.

Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding

4.1 Training details

We train on a combination of internal datasets, with $\approx 460\text{M}$ image-text pairs, and the publicly available Laion dataset [61], with $\approx 400\text{M}$ image-text pairs. There are limitations in our training data, and we refer the reader to Section 6 for details. See Appendix F for more implementation details.

6 Conclusions, Limitations and Societal Impact

Imagen's training data was drawn from several pre-existing datasets of image and English alt-text pairs. A subset of this data was filtered to removed noise and undesirable content, such as pornographic imagery and toxic language. However, a recent audit of one of our data sources, LAION-400M [61], uncovered a wide range of inappropriate content including pornographic imagery, racist slurs, and harmful social stereotypes [4]. This finding informs our assessment that Imagen is not suitable for public use at this time and also demonstrates the value of rigorous dataset audits and comprehensive dataset documentation (e.g. [23, 45]) in informing consequent decisions about the model's appropriate and safe use. Imagen also relies on text encoders trained on uncured web-scale data, and thus inherits the social biases and limitations of large language models [5, 3, 50].

It is in this backdrop we worryingly observe that the authors of the BASIC model paper have not even addressed model safety and dataset auditing issues in spite of having trained their model on the largest image-text dataset ever assembled and presented a full length 47 page paper with 3 revisions on ArXiv (See <https://arxiv.org/abs/2111.10050>).

E The ViT-L/14 OpenCLIP Network Architecture

The ViT-L/14 version of OpenCLIP has 428M parameters and 97M activations. Both the image and text branches output 768-dimensional embeddings. The image branch takes in images of size 224×224 , and has a depth of 24 layers and a hidden dimension of size 1024. On the other hand, the text branch has a depth of 12 layers and a hidden dimension of size 768.

The ViT-L/14 version of OpenCLIP was trained using 400M samples from the LAION-400M dataset with a batch size of 96 per GPU over 400 GPUs for a total batch size of 38400 for 32 epochs. The learning rate was set to 6×10^{-4} with 5000 warm-up iterations. The total training time for the model was 88 hours.

F On AllLookSameism, negative stereotypes and racial misclassification




Figure 13: Heatmap of the confusion matrix of the race-gender classification experiment showing misclassification Latino/Latina individuals as ‘Asian’ class. This misclassification got worse with dataset scaling.

The goal here was to understand how stereotypes about facial appearances are cross-related with racial identities. When we looked at the results (Figure 13) we saw an interesting theme emerge: That the self-identified Latino/Latina individuals were misclassified with high confidence as one of the ‘Asian’ classes on account of the presence of *epicanthic folds* and this tendency to stereotype got worse with dataset scaling. The title of these subplots here are formatted as strings with 4 fields separated by the ‘|’ character: $\langle \text{cfd_ViT-L-14} \rangle | \langle \text{training-dataset} \rangle | \langle P_{lf \rightarrow af} \rangle | \langle P_{lm \rightarrow am} \rangle$. Here, $P_{lf \rightarrow af}$ is the probability that an image belongs to the Latina-Female category was misclassified as Asian-Female (and $P_{lm \rightarrow am}$ is the probability that an image belong to the Latina-Male category was misclassified as Asian-Male). As seen in the first of the 3 subplots (from left) that maps to the OpenAI-WIT dataset 23 of the 56 latina women were misclassified as asian women leading to a $P_{lf \rightarrow af} = 23/56 = 0.411$. This misclassification rate was better for the LAION-400M model (0.125) and worsened to 0.179 for the LAION-2B-En model, thereby yielding yet another example of worsening of the bias-related metrics upon scaling the dataset from 400M to 2B samples. The same trend also showed up for Latino men with the misclassification rate increasingly nearly 50% from 0.077 to 0.115.

Correspondingly, there exists a substantial body of scientific literature (See [29, 38, 71, 68]) on not just the oft-ignored high levels of prevalence of the epicanthic folds in Hispanic/LatinX populations¹³ but also on the sociological ramifications of this *alllookslike-ism* [68] that permeates aspects of the mainstream culture.

¹³“In Latinos, the inner canthal distance and lateral canthal angle of inclination were similar to Asians, while the lid crease spanned the range from Asians to Caucasians. Half of the Latinos had epicanthal folds” [29]

G ‘Gorilla’ in the alt-text descriptions

In this appendix, we provide details pertaining to the six images constituting the collage in Figure 10. The URLs and the Alt-text descriptions have not been edited or sanitized and are presented verbatim as we found them in the dataset.

1. Sub-figure (a) was sourced from http://pigroll.com/img/youre_gonna_get_raped.jpg with the alt-text description: "Smiling albino gorilla wants to rape you"
2. Sub-figure (b) was sourced from https://content.wafflegirl.com/galleries/content/2/247/2247763_ec5249d.jpg with the alt-text description: "men fuck gorilla girl".
3. Sub-figure (c) was sourced from <https://media.istockphoto.com/photos/female-gorilla-with-baby-picture-id578309582> with the alt-text description: "Woman stripper with gorilla partner, girl fuck with eel"
4. Sub-figure (d) was sourced from https://new4.fjcdn.com/pictures/Gorilla+warfare_d41bd9_3543352.jpg with the alt-text description: "Gorilla warfare. . Top Comments This has been Bagged as more hide r Not Spam What the fuck did you just say about me, you little bitch? I' ll have you know I gr Gorilla warfare Top Comments This has been Bagged as more hide r Not Spam What the fuck did you just say about me little bitch? I' ll have know I gr"
5. Sub-figure (e) was sourced from https://farm1.static.flickr.com/552/31262527324_caa242e8d8_m.jpg with the alt-text description: "Mama Ayesha's Restaurant and Donald Trump (The Devils in the Details) Tags: donaldtrump mamaayeshas wallofpresidents hanksywashingtondc cia gop isis vladimirputin russia sexdrugsandrocksandroll hillaryclinton plannedparenthood bigot dumptrump thewalkingdead republican pedophile mikepence nastywoman badhombre conservative rape riencepriebus donalddmccahn stevenbannon frankgaffney jeffsessions generaljamesmattis generaljohnkelly stevenmnuchin andypuzder wilburross cathymcmorrisrodgers trumpforpresidentbobblehead poopydiaper ktmcfarland mikepompeo nikkihaley betsydevos tomprice scottpruitt seemaverma gorilla marriageequality kukluxklan daryldixon newyorkcity melaniatrump riggedelection jihad terrorist taliban mexicanwall racism confederateflag nazi islam freedom berniesanders americannaziparty thebeatles therollingstones democrat civilrights tednugent temptantrum contraception abortion tinfoilhatsociety michelleobama she'sacunt foxnews liberal"
6. Sub-figure (f) was sourced from <https://see.xxx/mt/sL/1994633.jpg> with the alt-text description: "Hirsute wet crack of this gorilla lady is so nasty that dont crave to fuck that".

