

Emergence of Text Readability in Vision Language Models

Jaeyoo Park^{1*} Sanghyuk Chun³ Wonjae Kim^{3†} Sangdoon Yun³ Bohyung Han^{1,2}

Computer Vision Laboratory, ECE¹ & IPAI², Seoul National University
Naver AI Lab³

{bellos1203, bhhan}@snu.ac.kr {sanghyuk.c, wonjae.kim, sangdoon.yun}@navercorp.com

Abstract

We investigate how the ability to recognize textual content within images emerges during the training of Vision-Language Models (VLMs). Our analysis reveals a critical phenomenon: the ability to read textual information in a given image (**text readability**) emerges abruptly after substantial training iterations, in contrast to semantic content understanding which develops gradually from the early stages of training. This delayed emergence may reflect how contrastive learning tends to initially prioritize general semantic understanding, with text-specific symbolic processing developing later. Interestingly, the ability to match images with rendered text develops even slower, indicating a deeper need for semantic integration. These findings highlight the need for tailored training strategies to accelerate robust text comprehension in VLMs, laying the groundwork for future research on optimizing multimodal learning.

1. Introduction

Vision-Language Models (VLMs) have significantly improved their ability to integrate visual and linguistic information [5–7, 14, 15, 22], achieving strong performance in diverse tasks. Specifically, CLIP [22] and its variants [6, 7, 28] demonstrate the ability to recognize objects, scenes, and texts within images [5, 20, 29], enabling applications such as scene-text recognition or document understanding [21, 31]. This progress highlights VLMs as powerful tools for bridging the gap between images and text, yet it also raises intriguing questions about the underlying mechanisms driving these capabilities.

Despite these advances, the training dynamics behind VLMs’ ability to process textual content within images remain underexplored. Recently, Lin et al. [17] identifies the source of such ability as images containing captions in

*Works done during an internship at NAVER AI Lab

†Now at TwelveLabs

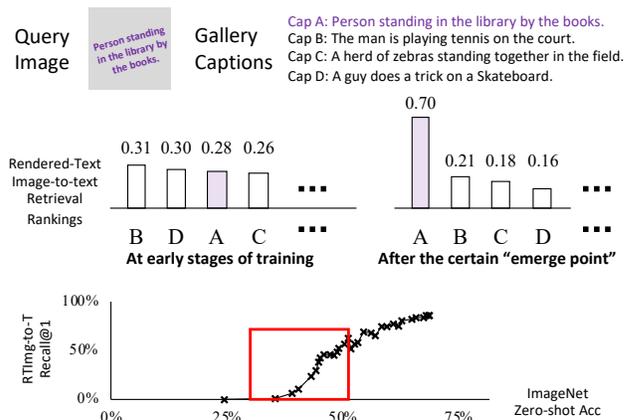


Figure 1. Emergence of *text readability* in VLMs during training. Our findings show that VLMs can suddenly learn to visually understand rendered text. Initially, the model struggles to recognize rendered text inputs (early-stage retrieval rankings, left). However, after a certain “*emerge point*”, its ability to interpret rendered text improves significantly (updated retrieval rankings, right). This sharp increase in Rendered-Text Image-to-Text (RTimg-to-T, bottom) occurs only after achieving a certain level of semantic understanding represented by ImageNet zero-shot accuracy.

large-scale web-crawled training data. However, Lin et al. mainly focus on the negative impact of images containing rendered text on semantic understanding of pure visual contents. This leaves open questions about the learning trajectories of how VLMs develop text readability and whether this development aligns with their semantic understanding.

In this paper, we focus on the *text readability* in VLMs, *e.g.*, can a VLM “read” the texts in an image? This can be measured by a retrieval task where a query is an image containing the text and the candidates are the corresponding text (*e.g.*, a rendered text caption such as the “Person standing in the library by the books” image and Cap A in Fig. 1). If a model can *visually* understand textual information in an image, the retrieval task can be easily solved. Our systematic study reveals an interesting insight regarding the text readability in VLMs: unlike semantic content under-

standing, which develops gradually from the initial stages, *the ability to read text emerges abruptly after substantial training*. This delayed emergence may reflect how optimizing contrastive loss prioritizes semantic learning, while text readability, requiring symbolic understanding, appears with further refinement. Fig. 1 summarizes our findings.

Moreover, we show that the ability to match images with rendered-text images emerges after additional training, lagging behind the development of rendered-text image-to-text (*RTImg-to-T*) matching. This underscores the complex interactions of visual and textual features, as the model initially relies on superficial pattern recognition before gradually acquiring deeper semantic integration. By highlighting this emergence, our work provides novel insights into the design and interpretability of multimodal systems.

The rest of the paper is organized as follows. Section 2 discusses related works on emergent phenomena of foundation models and text readability of VLMs. The details of our approach and experimental findings are described in Section 3. Finally, Section 4 concludes by summarizing our contributions and proposing future research directions.

2. Related Work

2.1. Emergent Abilities of Foundation Models

Recent studies have explored emergent capabilities in large-scale foundation models, encompassing both large language models (LLMs) and vision-language models (VLMs). Research has shown that LLMs exhibit abrupt improvements in tasks like complex arithmetic as their scale increases, suggesting a non-linear progression in ability [24, 27, 30]. For example, GPT models demonstrate proficiency in zero-shot translation, calculation, and action planning without specific training for these tasks [1, 3]. Recent findings challenge previous assumptions, suggesting that achieving low pre-training loss—rather than increasing model size or computational power—may be the critical factor behind these emergent abilities across various benchmarks [9].

In VLMs, a recent study observes that CLIP model acquires the ability to use a red circle as a visual prompt for region-specific focus, which is a skill absent in smaller models, showing how scaling reveals new functionalities [26]. This ability arises from web-crawled training data with red circles and matching captions. Our research reveals that text readability emerges as a distinct VLM capability, similarly shaped by training data with text in images.

2.2. Text Readability in VLMs

Recent studies have investigated the text readability of VLMs, uncovering a range of insights into their development. Wang et al. [29] first demonstrated that multimodal large language models, after large-scale pretraining, can read text embedded in images. Research into CLIP revealed



Raw Text T “Person standing in the library by the books”

Figure 2. (T, I, R) triplet example generated by FLUX.1 dev using the NoCaps dataset.

more intricate mechanisms. For instance, Materzyńska et al. [20] show that the image encoder of CLIP can separate word image spelling from natural scene meaning through projection layer finetuning. Additionally, Gandelsman et al. [12] identify specific attention heads specialized in recognizing letter patterns. However, Lin et al. [17] highlight a limitation known as “parrot bias,” where models merely learn to recognize the patterns of visible text in images, potentially missing deeper visual understanding. Our study explores the temporal dynamics of text-reading development in VLMs, examining its emergence during training to offer a fresh perspective on their evolution.

3. Method and Experimental Results

3.1. Dataset Construction

Our goal is to systematically investigate the emergence of text readability in VLMs. To achieve this, we require triplets consisting of (a) a text description T , (b) its corresponding natural image I , and (c) an image R with the visually rendered text of T – an example is shown in Fig. 2. In this subsection, we introduce how we construct these (T, I, R) triplets. Our comprehensive evaluation dataset comprises 3,000 high-quality triplets.

We generated these samples by applying a diffusion-based text-to-image pipeline to create images based on captions from the NoCaps dataset [2], which provides a diverse corpus of realistic image-caption pairs spanning multiple domains. For the diffusion model, we employed FLUX.1 dev [16], a recent diffusion model recognized for its exceptional text rendering fidelity. To ensure dataset diversity and robustness, we varied multiple rendering parameters, including font style, text color, background color, and spatial positioning of the text elements. This parametric variation allows for a more thorough evaluation of VLMs’ text comprehension ability across different visual contexts.

During dataset curation, we applied a rigorous quality control protocol to filter out samples where text was im-

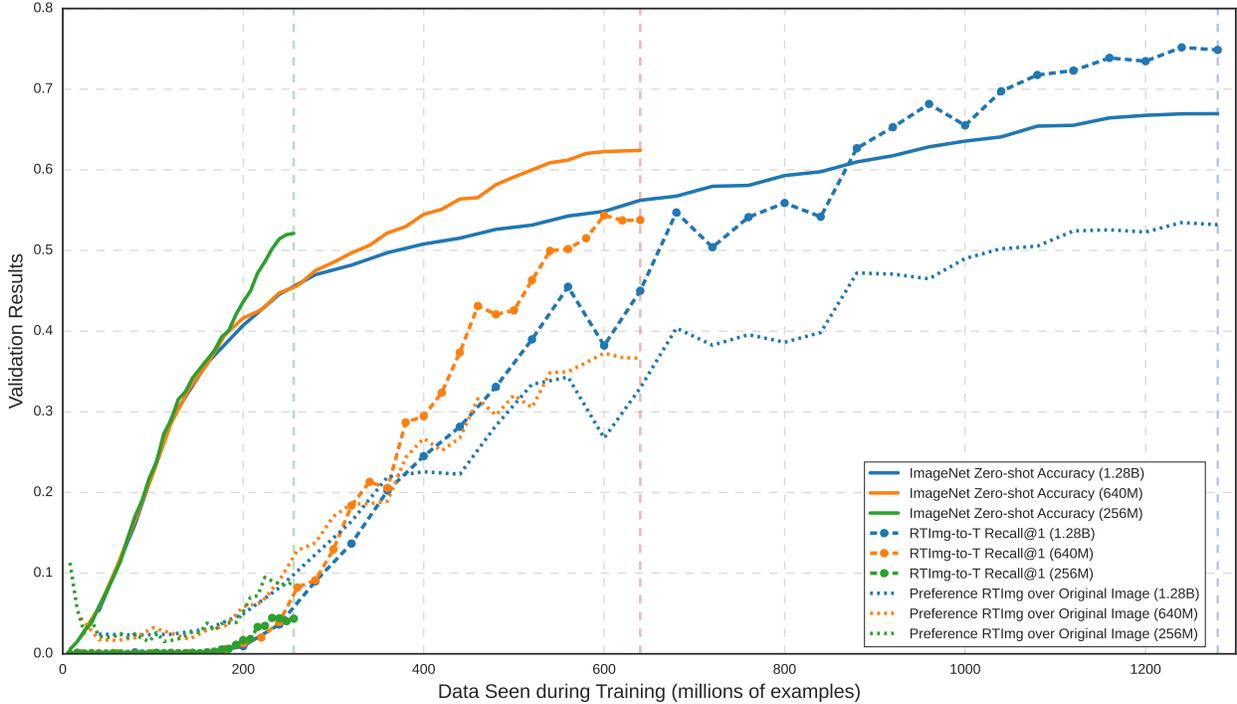


Figure 3. Validation performance of ViT-B/16-based CLIP models trained with three different sample sizes (256M, 640M, and 1.28B) as a function of data seen during training. Note that “RTImg” indicates “Rendered-Text Image”.

properly rendered or became illegible due to visual artifacts or poor contrast. This meticulous filtering process ensured that only high-fidelity examples remained in the final dataset, resulting in 3,000 validated triplets suitable for comprehensive analysis.

3.2. Evaluation Metrics and Protocol

We next explain how we quantitatively measure the text readability in VLMs. Our primary analysis uses a rendered-text image-to-text retrieval metric (RTImg-to-T Recall@1), which assesses the model’s ability to associate rendered-text images with their corresponding raw texts. Models with strong text recognition capabilities should achieve high retrieval performance on this task. Without this ability, its performance should be near random. Additionally, to assess deeper semantic understanding of rendered text, we employ a rendered-text image-to-image retrieval metric (RTImg-to-I Recall@1), detailed in Section 3.4. We also established a “preference” criterion, evaluating if the models assign higher similarity scores for rendered-text images over natural images when matched with the raw text.

We measure the rendered-text image-to-text retrieval metric and preference on ViT-B/16 CLIP models trained on the DataComp-1B dataset [11], a large-scale collection of 1.4 billion image-text pairs. We trained the models with various amounts of seen samples and measured the evaluation metrics during training. Notably, DataComp-1B likely

contains numerous image-caption pairs where captions are rendered directly within the images. This is primarily because it is constructed from web-crawled data, which naturally includes many rendered texts. Furthermore, as noted by Lin et al. [17], the DataComp filtering process may not effectively exclude “parroted captions”—captions that merely replicate text embedded in the visual content. Therefore, CLIP models trained on DataComp-1B are expected to eventually develop text readability. Our primary goal is to investigate *when* this ability emerges during training.

To investigate the impact of scale on model performance, our experiments utilize three distinct training dataset sizes (256M, 640M, and 1.28B seen samples). For each run, we checkpoint 32 intermediate weights with uniform intervals to measure the text readability during training. Each model is trained on DataComp-1B with batch size of 16,384 using the AdamW optimizer [19] by setting $\beta_1 = 0.9$ and $\beta_2 = 0.95$. During training, a random scaling (40% to 100%) and random color jittering are applied for data augmentation.

3.3. Emergence of Text Readability

To investigate how text readability develops during training, we analyzed the performance of ViT-B/16 CLIP models trained on DataComp-1B with varying sample sizes (256M, 640M, and 1.28B), as shown in Fig. 3. The figure plots ImageNet zero-shot classification accuracies (solid lines) and text readabilities (dashed lines), alongside preferences

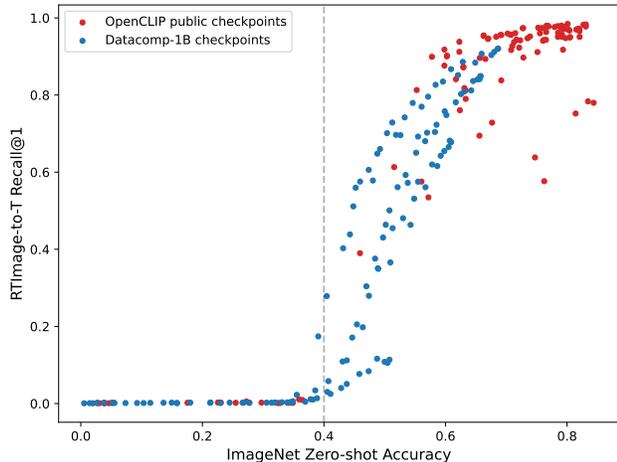


Figure 4. ImageNet zero-shot classification accuracy vs text readability for CLIP models. OpenCLIP public checkpoints (red) and Datacomp-1B checkpoints (blue) are distinguished by color. The plot highlights a non-linear correlation, with text readability emerging around 0.4 ImageNet zero-shot accuracy.

(dotted lines) of each model. Notably, text readability and the preference for rendered text images emerge consistently around 200 million samples, irrespective of the learning schedule or sample size. Despite the abundant presence of rendered captions in the training dataset and consistent improvements in ImageNet zero-shot accuracy, text readability emerges later in training. This delayed emergence may reflect how contrastive loss optimization tends to prioritize semantic learning, while symbolic representation learning—linked to text readability—appears to develop with further loss optimization.

Building on our previous analysis, we further investigated the relationship between text readability and semantic understanding by evaluating 114 publicly available CLIP models from the OpenCLIP repository [13]. For these public models, alongside our trained models, we assessed text readability and its correlation with semantic understanding. Semantic understanding was measured using ImageNet zero-shot classification accuracy [23], a standard metric for CLIP performance. Note that the public OpenCLIP weights include diverse architectures (*e.g.*, ViT-G [8], ConvNext [18], SigLIP [32]) and pretraining datasets (*e.g.*, LAION-2B [25], WebLi [4], DFN [10]), allowing for a broader examination.

Fig. 4 illustrates this relationship, distinguishing OpenCLIP models (red) from the Datacomp-1B-trained checkpoints (blue) by color. From the results, we observed a non-linear correlation between ImageNet zero-shot accuracy and text readability, with readability emerging abruptly around an accuracy of 0.4 regardless of the model type. This pattern holds consistently for Datacomp-1B-trained models across different seen samples during training (256M, 640M,

1.28B), indicating that the abrupt emergence of text readability is generalized to variations in training scale. This observation aligns with our hypothesis that the contrastive loss optimization process prioritizes semantic content understanding over symbolic representation learning, though both capabilities eventually develop with sufficient training.

3.4. Semantic Understanding of Rendered Text

Thus far, our analysis has focused on RTImage-to-T matching, which serves as an indicator of the model’s text readability. However, this approach may reflect a naïve pattern-matching strategy as Lin et al. [17] claimed, raising the question of whether the model can truly grasp the semantic meaning of rendered text. If a model truly understands rendered text and accurately embeds the meaning of the text in its representation space, the visual embedding of rendered text image (R), the textual embedding of the original caption (T), and the embedding of its corresponding natural image (I) should be closely aligned. Namely, we expect that not only the RTImage-to-T ability emerges but also the RTImage-to-I ability emerges similarly.

We report the RTImage-to-I Recall@1 in Fig. 5 along with RTImage-to-T Recall@1. This comparison reveals two key findings. First, RTImage-to-I recall is significantly lower compared to RTImage-to-T recall (*e.g.*, below 10% vs. above 70%). This finding may suggest that the emergence of the RTImage-to-T ability could be a reliance on superficial pattern recognition rather than deep semantic understanding.

Second, the RTImage-to-I matching ability emerges later than the emergence of the RTImage-to-T matching ability. We presume that this delay is likely because contrastive learning primarily involves direct comparisons between images and texts, rather than between rendered-text images and images—despite both being visual inputs. However, this RTImage-to-I capability gradually improves with sufficient training (*e.g.*, achieving around 10% Recall@1 with over 1B seen samples). Furthermore, we also observe that if we use a stronger model (*e.g.*, ViT-L/16), the RTImage-to-I recall@1 is substantially improved (*e.g.*, 20%). More details are in the next subsection.

Overall, we presume that the emergence of RTImage-to-T Recall@1 would not solely originate from the truly semantic and visual understanding of rendered text, but it also could rely on a simple pattern matching of rendered text (*e.g.*, detecting a specific word, rather than understanding the whole caption semantically). Developing a training strategy to achieve both matching abilities (RTImage-to-T and RTImage-to-I) will be an interesting future research direction.

3.5. Scaling Model Capacity with ViT-L/16

We extended our analysis to a larger ViT-L/16 backbone, training a CLIP model on the DataComp-1B dataset. This ViT-L/16 model used the same training configuration as

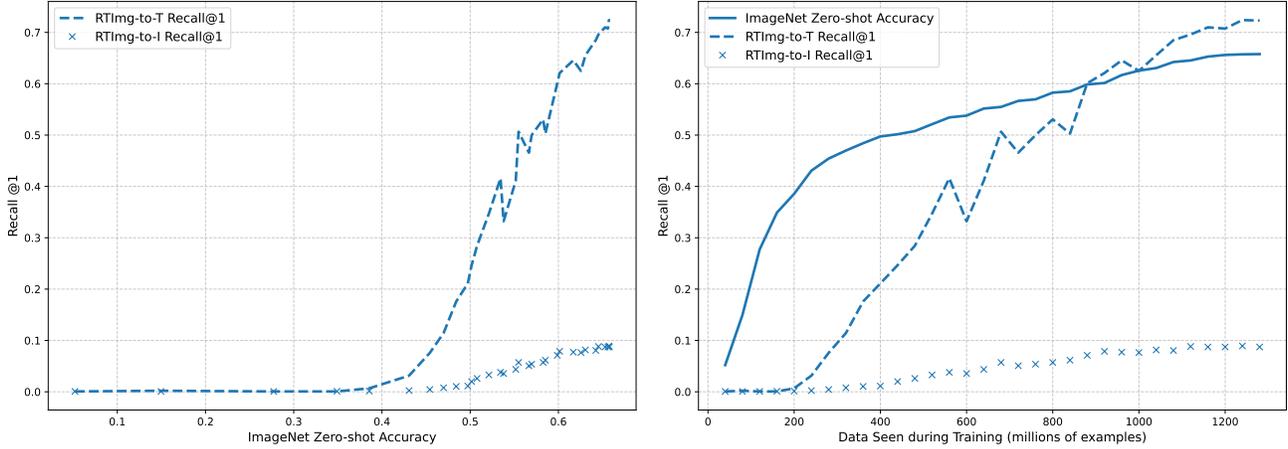


Figure 5. Comparison of CLIP models’ matching performance between rendered-text images and text versus rendered-text images and natural images across learning trajectory. The graph shows lower performance in RTImg-to-I matching, indicating reliance on pattern recognition, and delayed emergence of this ability compared to RTImg-to-T matching, reflecting the challenges of contrastive learning in aligning visually processed inputs.

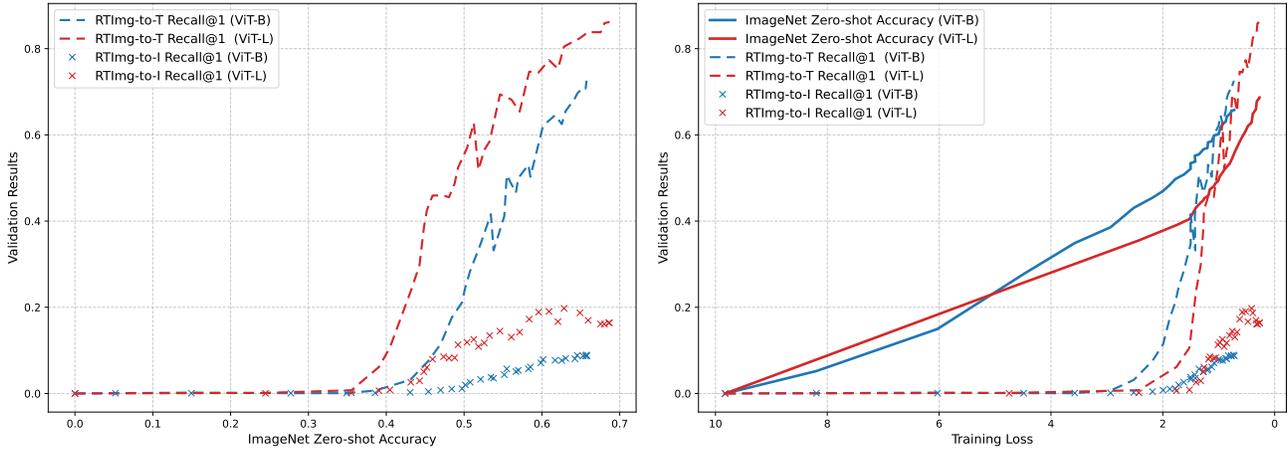


Figure 6. Comparison of text readability and semantic understanding of rendered-text images for ViT-B/16 and ViT-L/16 models, plotted against ImageNet zero shot accuracy and training loss.

its ViT-B/16 counterpart, trained with 1.28B seen samples. Following the methodology outlined by Du et al. [9], we also plotted the text readability and semantic understanding of rendered-text images for both ViT-B/16 and ViT-L/16 against training loss in Fig. 6. The results indicate that the pattern previously observed in the smaller model—where text readability and semantic understanding of rendered text emerge after the model grasps semantic understanding of images—also holds for the larger ViT-L/16; and furthermore, a larger model exhibits stronger text readability.

4. Conclusion

In this paper, we have explored the emergence of textual content recognition in Vision-Language Models (VLMs), uncovering significant patterns in their training dynamics.

Our central finding is that the ability to read text (RTImg-to-T) arises abruptly after extensive training, a stark contrast to the steady development of general semantic understanding. This delay likely reflects how contrastive learning prioritizes broader visual-semantic features before symbolic capacities such as text recognition. Models seem to first prioritize general visual-semantic features before developing symbolic capacities like text recognition. The difficulty of tasks like RTImg-to-I further indicates that deeply integrating rendered text visually is a more advanced capability. These insights emphasize the importance of tailoring training approaches to foster not just text recognition but genuine understanding in VLMs. We hope this study sets the stage for future investigations into refining multimodal learning strategies.

Acknowledgement This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) [No.RS-2021-II211343, Artificial Intelligence Graduate School Program (Seoul National University)].

References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 2
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. nocaps: novel object captioning at scale. In *ICCV*, 2019. 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020. 2
- [4] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, et al. Pali: A jointly-scaled multilingual language-image model. In *ICLR*, 2023. 4
- [5] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *CVPR*, 2024. 1
- [6] Sanghyuk Chun, Wonjae Kim, Song Park, and Sangdoon Yun. Probabilistic language-image pre-training. In *ICLR*, 2025. 1
- [7] Karan Desai, Maximilian Nickel, Tanmay Rajpurohit, Justin Johnson, and Shanmukha Ramakrishna Vedantam. Hyperbolic image-text representations. In *ICML*, pages 7694–7731. PMLR, 2023. 1
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 4
- [9] Zhengxiao Du, Aohan Zeng, Yuxiao Dong, and Jie Tang. Understanding emergent abilities of language models from the loss perspective. In *NeurIPS*, 2024. 2, 5
- [10] Alex Fang, Albin Madappally Jose, Amit Jain, Ludwig Schmidt, Alexander Toshev, and Vaishal Shankar. Data filtering networks. In *ICLR*, 2024. 4
- [11] Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruva Ghosh, Jieyu Zhang, et al. Datacomp: In search of the next generation of multimodal datasets. In *NeurIPS*, 2023. 3
- [12] Yossi Gandelsman, Alexei A Efros, and Jacob Steinhardt. Interpreting clip’s image representation via text-based decomposition. In *ICLR*, 2024. 2
- [13] Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishal Shankar, Hongseok Namkoong, John Miller, Hananeh Hajishirzi, et al. Openclip, 2021. 4
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 1
- [15] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *ICML*, 2021. 1
- [16] Black Forest Labs. Flux. <https://github.com/black-forest-labs/flux>, 2024. 2
- [17] Yiqi Lin, Conghui He, Alex Jinpeng Wang, Bin Wang, Weijia Li, and Mike Zheng Shou. Parrot captions teach clip to spot text. In *ECCV*, 2024. 1, 2, 3, 4
- [18] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 4
- [19] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019. 3
- [20] Joanna Materzyńska, Antonio Torralba, and David Bau. Disentangling visual and written concepts in clip. In *CVPR*, 2022. 1, 2
- [21] Jaeyoo Park, Jin Young Choi, Jeonghyung Park, and Bohyung Han. Hierarchical visual feature aggregation for ocr-free document understanding. In *NeurIPS*, 2024. 1
- [22] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 1
- [23] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. In *IJCV*, 2015. 4
- [24] Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are emergent abilities of large language models a mirage? In *NeurIPS*, 2023. 2
- [25] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. 4
- [26] Aleksandar Shtedritski, Christian Rupprecht, and Andrea Vedaldi. What does clip know about a red circle? visual prompt engineering for vlms. In *ICCV*, 2023. 2
- [27] Charlie Snell, Eric Wallace, Dan Klein, and Sergey Levine. Predicting emergent capabilities by finetuning. In *COLM*, 2024. 2
- [28] Michael Tschannen, Basil Mustafa, and Neil Houlsby. Clippo: Image-and-language understanding from pixels only. In *CVPR*, 2023. 1
- [29] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. GIT: A generative image-to-text transformer for vision and language. *TMLR*, 2022. 1, 2
- [30] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *TMLR*, 2022. 2

- [31] Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, et al. UReader: universal ocr-free visually-situated language understanding with multimodal large language model. In *EMNLP*, 2023. [1](#)
- [32] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *ICCV*, 2023. [4](#)