

A Generative Approach for Wikipedia-Scale Visual Entity Recognition

Mathilde Caron Ahmet Iscen Alireza Fathi Cordelia Schmid
 Google Research

Abstract

In this paper, we address web-scale visual entity recognition, specifically the task of mapping a given query image to one of the 6 million existing entities in Wikipedia. One way of approaching a problem of such scale is using dual-encoder models (e.g. CLIP), where all the entity names and query images are embedded into a unified space, paving the way for an approximate kNN search. Alternatively, it is also possible to re-purpose a captioning model to directly generate the entity names for a given image. In contrast, we introduce a novel Generative Entity Recognition (GER) framework, which given an input image learns to auto-regressively decode a semantic and discriminative “code” identifying the target entity. Our experiments demonstrate the efficacy of this GER paradigm, showcasing state-of-the-art performance on the challenging OVEN benchmark. GER surpasses strong captioning, dual-encoder, visual matching and hierarchical classification baselines, affirming its advantage in tackling the complexities of web-scale recognition.

1. Introduction

Generative vision-language models such as GPT-4 [30], Flamingo [2] or PALI [5], are becoming increasingly popular for computer vision applications. They show an impressive ability to generate free-form text for describing the contents of an image (captioning), or answering questions based on an image (visual-question answering). Nevertheless, their potential for *recognition* tasks [12], which usually require a more concise, structured output, remains under-explored. The focus of this paper is to explore their application for the challenging task of web-scale entity recognition. A recent benchmark, Open-domain Visual Entity recognition (OVEN) [12], challenges models to associate an image with a Wikipedia entity from a pool of over six million entities. Models must establish a robust association between images across millions of coarse-grained and fine-grained entities, encompassing a wide spectrum of concepts such as animals, buildings, locations, and a multitude of others [12].

Code: github.com/google-research/scenic/tree/main/scenic/projects/gerald

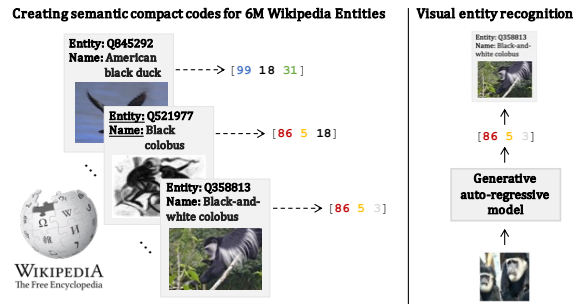


Figure 1. We introduce GER, a novel generative paradigm for web-scale visual entity recognition. We create compact semantic codes for each entity, and learn to auto-regressively generate them for a given query image at inference.

Traditionally, the predominant methods employed to address the challenge of visual entity recognition have revolved around either classification or contrastive dual-encoder paradigm like CLIP [32]. While classification offers a straightforward approach, it grapples with limitations when confronted with extensive label spaces such as that of OVEN, resulting in substantial parameter counts and practical engineering complexities. The dual-encoder approach on the other hand, learns a unified image-text feature space, thereby facilitating efficient nearest neighbor searches for recognition. Nonetheless, this approach exhibits its own drawbacks: (a) it does not directly optimize for the final recognition task but instead relies on indirect optimization through contrastive loss where a set of negative data has to be subsampled at training time [11, 29, 32], (b) compressing either the image or text into an embedding vector results in loss of information, detrimentally affecting performance for fine-grained recognition [15] and (c) the memory requirements for storing dense representations scale proportionally with the size of the entity set.

These challenges of the dual-encoder paradigm have kindled interest in alternative strategies. Notably, in Natural Language Processing (NLP) domain, recent works challenge the dual-encoder approach and use generative models instead for information retrieval [6, 25, 31, 33, 41, 42]. These works represent each element of the corpus by a compact code of integers, and learn an auto-regressive generative model to decode the target code for a given query. This

paradigm promises to overcome some drawbacks of dual-encoders by simplifying the retrieval pipeline such that the training and inference objectives are the same, and directly encoding the corpus within the model’s parameters. Also as an alternative to dual encoders, OVEN paper [12] showcases the feasibility of extending a generative image captioning model [5] for visual entity recognition by matching the generated caption to one of the Wikipedia entity texts [34].

Inspired by these recent explorations, we propose a Generative Entity Recognition (GER) framework (illustrated in Fig. 1) to facilitate end-to-end visual entity recognition by leveraging generative auto-regressive models. Specifically, we represent each Wikipedia entity with a *code*, *i.e.* a short sequence of integers. Then, we train models to predict an entity from an input image by auto-regressively generating the code corresponding to the target entity. We find that creating unAmbiguous, Language-based and Discriminative (ALD) entity codes results in the best variant of our GER framework, which we denote by GER-ALD. In fact, while we observe that unstructured “atomic” codes work well in some scenarios, they fail when training data or model capacity are limited or more importantly, when the entity set reaches the million scale (see Sec. 4.4.1). Plus, they cannot generalize to new entities. In contrast, we find that semantically-structured codes based on language improve upon atomic codes by leveraging generic concepts shared across related entities (see example in Fig. 1 with “Black colobus” and “Black-and-white colobus” sharing common code tokens). A simple way of creating codes based on language is to directly tokenize [20] the entity name, which is akin to image captioning where the entity name is used as a caption [6, 12]. However, we find that such tokenized entity names contain clutter and noisy information, all the more so when the entity name is long (see Sec. 4.4.2). Our GER-ALD method improves over this simple captioning baseline by decoding only the most *discriminative* part of the tokenized entity name, *i.e.* the part which makes the considered entity name the most different compared to all other entities.

Finally, we also propose an entity-based pre-training to condition the GER models to web-scale entity recognition. Inspired by recent advances in retrieval-based methods [15, 23], we retrieve a subset of images from a large-scale image-text dataset typically used for captioning or contrastive pre-training [5] and re-purpose it by replacing the original text captions with related OVEN entity names. Overall, our experiments demonstrate the efficacy of the proposed GER paradigm: GER-ALD outperforms previously published numbers on OVEN benchmark [12] by +6.7 top-1 accuracy, while using 42× less parameters. In summary, our contributions are as follows:

- a generative entity recognition framework (GER) to facilitate end-to-end visual entity recognition;
- an innovative strategy for encoding Wikipedia enti-

ties into unambiguous language-based discriminative (ALD) codes that are highly effective for GER;

- an entity-based pre-training process without requiring human intervention;
- state-of-the-art results in challenging web-scale OVEN entity recognition and on-par performance to traditional classifiers in smaller-scale label-space scenarios.

2. Related work

Visual entity recognition aims to recognize classes, or entities given visual inputs [35]. Granularity of visual entity recognition tasks varies from every-day generic objects [8, 9], to fine-grained domains, such as birds [44], dogs [17], cars [18], food [4], landmarks [47], faces [50] and natural world species [43]. Some challenges for the visual entity recognition tasks include imbalanced training classes following a long-tailed distribution [24], or noisy training labels [22]. Recent work [12] proposes a new, *web-scale* dataset for open-domain entity recognition. This challenging benchmark contains 6M entity names derived from Wikipedia page titles, including coarse-grained and fine-grained entities, encompassing a wide spectrum of concepts such as animals, buildings, organizations, landmarks, and a multitude of other. The authors show that generative captioning models (*i.e.* PaLI [5]) outperform dual encoder models for large-scale entity recognition. In this paper, we build upon this observation, and study generative models for accurate and efficient entity recognition.

Extreme classification tackles entity recognition specifically at a very large scale with a pure classification approach [1, 3, 26]. Typical approaches explore strategies for scaling to the hundred of thousands scale and preliminary results are even shown at million scale [1]. By leveraging generative image-to-text models, we propose a fresh perspective beyond traditional classification methods typically used in the context of large-scale visual entity recognition.

Generative auto-regressive retrieval methods are increasingly popular in NLP [6, 25, 31, 33, 41, 42]. GENRE retrieves Wikipedia entities by generating their names in an autoregressive fashion. Seminal work DSI [42] shows the benefit of learning to decode compact codes (created either randomly or with hierarchical k-means clustering) associated with each document. Neural Corpus Indexer [46] proposes a specific decoding scheme for generative retrieval and show the benefit of *query augmentation* by automatically generating training queries for documents to be indexed. TIGER [33] studies generative retrieval in the context of recommender systems. Finally, [31] conducts a systematic study of generative retrieval systems when scaled to millions of document passages. Only very few works explore this family of approaches in computer vision domain, and only in very small-scale and uni-modal scenarios [49].

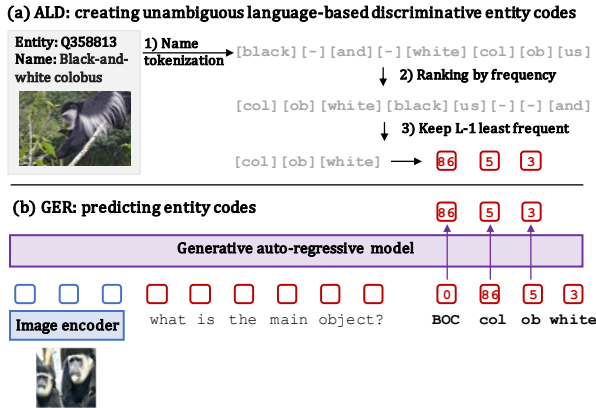


Figure 2. **Overview of GER-ALD method.** (a) We utilize a text tokenizer to create compact and semantic codes, which represents each entity with short, but discriminative representations. (b) We learn a generative auto-regressive model, which learns to decode the correct code for given query image and text pair.

3. Method

Our goal is to explore how to adapt generative auto-regressive models to the task of visual entity recognition (GER). While previous works have shown preliminary signal that it is possible to repurpose autoregressive models for entity recognition by directly decoding entity names [6, 12], we propose a more effective strategy. An overview of our framework is in Fig. 2.

3.1. Problem definition

Web-scale visual entity recognition. The Open-domain Visual Entity recognition (OVEN) [12] task consists of mapping input visual queries to one of the 6M English Wikipedia entities. More specifically, for a given image query x_v and text query x_t , the model needs to recognize the corresponding entity e among the set \mathcal{E} of all possible entities. The purpose of the input text x_t is to achieve unambiguous recognition. For example, when several entities are represented in the query image x_v , the text query indicates which one needs to be recognized. Each entity $e \in \mathcal{E}$ comes with an entity name, denoted by t_e , which corresponds to the title of the entity Wikipedia page.

Representing each entity with a code. In GER, we represent each entity e by a *code* denoted by $c^e = \{c_1^e, \dots, c_L^e\} \in \llbracket 1, V \rrbracket^L$ where L is the length of the code and V is the size of the *vocabulary* of all integer values that each code token c_i^e can take. This forms up to V^L unique codes. Note that vanilla image classification and captioning baselines can both be cast into this code formulation. In fact, with $L = 1$ and $V = |\mathcal{E}|$, the codes are equivalent to the labels used in standard multi-class classification. On the other hand, if each code token value in $\llbracket 1, V \rrbracket$ maps to a (sub-)word in a pre-defined vocabulary [20], then the codes simply correspond to standard tokenized text used in captioning

models [19, 39, 45]. In the following paragraphs, we detail GER-ALD, our most effective strategy for building codes C to represent all 6M English Wikipedia entities.

3.2. GER-ALD: Creating ALD codes for GER

We design the code set C so that it has three properties which we find are important for effective GER models: i) semantically structured thanks to language, ii) discriminative and compact, and iii) unambiguous. Our algorithm to create such unambiguous, language-based and discriminative codes, called ALD, is illustrated in Fig. 2 (a) and described in pseudo-code in Algorithm 1 of the Appendix.

Semantic tokens based on language. We find that entity codes C benefit from following a semantic structure, especially in scenarios where memorizing unstructured atomic codes is difficult. We show in Sec. 4.4.1 that using unstructured atomic codes fail when the amount of training data or the model capacity are limited or, of particular interest, when the entity set size increases to the million scale (see Fig. 3). Intuitively, we want entities that are semantically similar to have some overlapping code tokens. For example, we wish that entities $e = Q521977$ with corresponding name $t_{Q521977} = \text{“Black colobus”}$ and $e = Q358813$ with corresponding name $t_{Q358813} = \text{“Black-and-white colobos”}$ to share some code tokens, given that these correspond to two close species.

A simple yet effective way of having semantic codes is to tokenize the entity names based on text tokenizers [6, 19, 20, 39]. If each of the sub-words in the entity names are mapped to an integer representing this sub-word, then entities Q358813 and Q521977 naturally share code tokens: those representing the phrase “colobus”. We denote by $\Phi(\cdot)$ an off-the-shelf text tokenizer with a vocabulary of V_Φ sub-words such that $\Phi(t_e) = \{y_1^e, \dots, y_{L_e}^e\} \in \llbracket 1, V_\Phi \rrbracket^{L_e}$ where L_e is the length of the tokenized entity name $\Phi(t_e)$. In practice we use the same language tokenizer as GIT [45] for $\Phi(\cdot)$ and have a vocabulary size of $V = V_\Phi = 30522$. We refer to the baseline of using codes C created by simple tokenization of the entity name as GER-CAPTION (i.e. we treat the entity name as a caption) [6]. We show in the following paragraph how GER-ALD codes differ from such GER-CAPTION codes by making them more compact and discriminative.

Discriminative and compact codes. Our goal is to build short and highly discriminative codes because they are easier to learn for the model, as validated by our experiments in Sec. 4.4.2. For example, the tokenized entity name $\Phi(t_{Q358813}) = \Phi(\text{“Black-and-white colobus”})$ counts $L_{Q358813} = 8$ tokens, but clearly not all 8 tokens are important to make this entity discriminative compared to all other existing entities. Hence, we choose to represent each entity with the *bare minimum*, removing all the *clutter* which is not only non-discriminative but also adds noise. We achieve this by selecting the most discriminative and rarest tokens

within the tokenized entity name. Specifically, we compute the frequency f_v of each token value $v \in [1, V]$ in the vocabulary over the entire corpus of tokenized entity names $\{\Phi(t_e)\}_{e \in \mathcal{E}}$. We have $f_v = \frac{n_v}{\sum_{u=1}^V n_u}$ where n_v is the number of times v appears in $\{\Phi(t_e)\}_{e \in \mathcal{E}}$. We create an ALD code c_e for each entity by keeping only the $(L - 1)$ tokens with the lowest frequencies and discarding the other ones. For example for entity Q358813, the 3 tokens with the lowest frequencies are “col”, “ob” and “white”. Interestingly, these 3 most discriminative tokens appear at the end of the code for GER-CAPTION. By contrast, they appear right at the beginning of the code for GER-ALD and they constitute the only tokens to be decoded by the model, which intuitively explains the improved performance of GER-ALD codes, as analyzed later in Sec. 4.4.2 especially when entities have long names (see Fig. 4). Finally an interesting by-product of using short codes is that they are faster to decode (the complexity of decoding is $\mathcal{O}(L^2)$) and require less memory footprint to store.

Unambiguous codes. Note that several entities might share the same least frequent $(L - 1)^{\text{th}}$ tokens. In this case their code are exactly identical up to the $(L - 1)^{\text{th}}$ token. We use the last L^{th} token to ensure that each entity has a unique code: we greedily assign the last code token c_L^e to the next least frequent word of the tokenized entity name until the code c_e is different from all existing codes. If this still fails to create a unique code, we assign c_L^e to a random token value v' so that the resulting code is unique. With code length $L = 4$, only 0.5% of the entities use a random token value.

3.3. Training

In this section, we describe the model used to decode entity codes from an input image-text pair. Importantly, we also introduce our entity-based pre-training to condition the generative model to the task of entity recognition.

Auto-regressive generative models. We build upon GIT [45], an auto-regressive image-to-text generative model. The query image-text pair (x_v, x_t) is transformed into a set of d -dimensional embeddings using a visual encoder for x_v and the text tokenizer $\Phi(\cdot)$ for x_t . The resulting output is represented by $\mathbf{X}_v \in \mathbb{R}^{N_v \times d}$ (resp. $\mathbf{X}_t \in \mathbb{R}^{N_t \times d}$) for image (resp. text) tokens. We then input \mathbf{X}_v and \mathbf{X}_t to a decoder network $g(\cdot)$ whose task is to decode the next code token c_i^e , conditioned on the previous tokens $c_{j < i}^e$. Each code token value v in $[1, V]$ maps to a learnable d -dimensional vector \mathbf{Y}_v (gathered in the embedding matrix $\mathbf{Y} \in \mathbb{R}^{(V+1) \times d}$ where \mathbf{Y}_0 corresponds to the “beginning of code” token). We train with a language modeling loss:

$$\mathcal{L}^e = \frac{1}{L} \sum_{i=1}^L \ell(c_i^e, g([\mathbf{X}_v; \mathbf{X}_t; \mathbf{Y}_0; \mathbf{Y}_{c_{0 < j < i}^e}]))$$

where $[\cdot; \cdot]$ corresponds to the concatenation operation in the first dimension and ℓ is the softmax cross-entropy loss with label-smoothing [27]. We average \mathcal{L}^e over a mini-batch and

learn the weights of the visual encoder, decoder $g(\cdot)$ and embedding matrix \mathbf{Y} through back-propagation. When decoding, we use beam search to obtain the best predicted entity coded. We find that we do not need to constrain the beam search to existing codes since more than 99% of the top-1 predictions are valid codes for converged GER models.

Entity-based pre-training. Common auto-regressive models such as GIT [45] or PaLI [5] are pre-trained for descriptive captioning. As shown in Tab. 5 and Fig 9 of the Appendix, they generalize poorly to entity recognition. This is because of the task discrepancy between predicting a descriptive caption and predicting an entity name. In order to condition our models better for entity recognition, we propose to collect a significant number of entity-based pretraining images, each associated with a Wikipedia entity instead of a generic caption. However, such an entity-based pre-training dataset does not exist. We create it in an automatic way, without any human supervision.

To do so, we leverage existing large-scale image-caption datasets [37, 38]: unless specified otherwise we use WebLI [5]. For each Wikipedia entity, we retrieve in WebLI the image-caption pairs that best represent this entity and replace their original captions by this entity name [15, 23]. Specifically, we embed the 6M entity names of OVEN with a semantic text encoder [32] and find the top- k most similar captions in WebLI. We retrieve their corresponding images and replace their original captions by the considered entity name. We ensure that no image is assigned to multiple entities to avoid instability during training. We vary the number of retrieved images k per entity from 2 to 100 to produce pre-training datasets of different sizes: from 11M up to 55M images (see Fig. 6). We denote by Entity-WebLI (resp. Entity-LAION) the resulting dataset used for entity-based pretraining, built from WebLI (resp. LAION [38]). This way of creating pre-training data is akin to the query generation techniques used for generative retrieval in NLP [46]. However, rather than generating a synthetic input, we simply retrieve input images from a large-scale dataset.

3.4. Baselines

We compare our method to the following different baselines.

Hierarchical classification. Solving million-scale entity recognition with classification is unpractical due to the very large number of classes. A workaround is to use *hierarchical* classifiers. As OVEN does not come with hierarchical labels we obtain a 3-level hierarchy through k-means of the 6M entity names encoded with sentence-T5 [28]. We train a multi-class classifier for each parent node in the hierarchy. To avoid training a huge number of different classification matrices, we learn a generic classifier matrix per level which is modified by learnable small *modifiers* depending on the path in the hierarchy.

Dual encoders. Another typical workaround to classifica-

tion is to rely on deep metric learning approaches [36] such as Noise Contrastive Estimation [11] and its InfoNCE variant [29] as used in popular dual encoder approaches [16, 32]. Dual encoders learn a unified image-text feature space with separate encoders, thereby facilitating efficient nearest neighbor searches for recognition. We use CLIP-L/14 [32].

Visual matching. We also experiment with pure visual matching baselines. We use off-the-shelf CLIP-L/14 visual encoder and Entity-WebLI (55M) dataset as the memory. We use $k = 500$ for nearest neighbor search with majority voting as it obtains the best results on OVEN val set.

Captioning. We compare to Git-Large [45] or PaLI [5] image-to-text auto-regressive captioning models.

GER-baselines: alternative code creation strategies. We compare GER-ALD, *i.e.* the best variant of GER, with several alternatives. First, GER-ATOMIC refers to using atomic, completely unstructured codes, *i.e.* each code token c_i^e is randomly drawn from $[1, V]^L$ [42]. Second, we consider two alternatives using semantically structured codes: (i) GER-HKC where we embed the entity names with a pre-trained text encoder before applying hierarchical k-means clustering on the resulting embeddings [42] and (ii) GER-CAPTION where we create a code by tokenizing the entity name with $\Phi(\cdot)$ [6, 12]. Details on the baselines are in Appendix Sec. 6.4.

4. Experiments

In this section, we detail our experimental setup, compare our method with state of the art and baselines, and finally present thorough analyses on code creation and pretraining.

4.1. Experimental setting

OVEN dataset consists of 6,063,945 different entities [12]. We evaluate the models on the validation and test splits, by reporting the harmonic mean (HM) of top-1 accuracy scores between “seen” and “unseen” entities. Seen are entities present in the OVEN training set. Unseen entities are a subset of entities among the ones not present in the training set. The models are evaluated on a total of 3192 entities (1721 for seen and 1471 for unseen) for validation and 15888 entities (8355 for seen and 7533 for unseen) for test. We call the entities that the model is evaluated on by “positive” entities (*i.e.* the union of the 3192 validation and 15888 test entities) and all other entities by “negative” entities.

Pretraining and finetuning. Unless specified otherwise, we pretrain our models on the entity-WebLI dataset, which we create considering *all* 6M entity names as described in Sec. 3.3. After this entity-based pretraining, the models are finetuned on OVEN training set which consists only of the “seen” entities. All implementation details are in Sec. 6 in Appendix and code is released in the SCENIC library [7].

Preventing data leakage. We remove pretraining images from Entity-WebLI and Entity-LAION with a cosine simi-

Method	#par.(B)	Pretraining		OVEN test		
		dataset	#imgs	HM	seen	unseen
<i>Dual encoder approaches</i>						
CLIP _{VIT-L14}	0.42	OpenAI	400M	5.2	5.6	4.9
CLIPfusion _{VIT-L14}	0.88	OpenAI	400M	8.4	33.6	4.8
CLIP2CLIP _{VIT-L14}	0.86	OpenAI	400M	11.5	12.6	10.5
<i>Captioning approaches</i>						
GiT-Large	0.40	WebLI	100M	7.0	17.6	4.3
PaLI-3B	3	WebLI	1B	9.1	19.1	6.0
PaLI-17B	17	WebLI	1B	16.0	28.3	11.2
<i>Generative entity recognition</i>						
GER-ALD [‡] (Ours)	0.40	Entity-LAION	41M	20.9	29.1	16.3
GER-ALD (Ours)	0.40	Entity-WebLI	55M	22.7	31.5	17.7

Table 1. **Comparison with state-of-the-art** approaches on OVEN entity test split. We report the harmonic mean (HM) of the seen and unseen splits (top-1 accuracy) after finetuning on OVEN training set. Numbers are taken from [12] except methods based on GiT-Large which are run by us. We indicate the total number of parameters of each model (“# par.”) in billion and the pretraining dataset details. ‡: use only publicly available data.

ilarity (with CLIP-L/14 visual features) above 0.95 with any of the OVEN test or val images. We chose a 0.95 conservative threshold by looking at some examples: similarity 0.95 corresponds to conceptually similar images but clearly not duplicates (see Fig. 8 in Appendix).

4.2. Comparison with the state of the art

In Tab. 1, we compare the performance of GER-ALD, our best GER variant, on the OVEN entity benchmark with previously published numbers after finetuning on the OVEN training set. We see that our method outperforms previously proposed approaches by significant margins. Notably, GER-ALD improves over the captioning model PALI-17B by +6.8 top-1 HM test accuracy (a relative improvement of 43%) while using 42× less parameters.

4.3. Comparison with baselines

In Tab. 2, we compare GER-ALD with the different baselines described in Sec. 3.4. All baselines use exactly the *same pretraining dataset* entity-based WebLI (55M) and model architectures of comparable sizes.

Comparing GER to different paradigms. We see in Tab. 2 that GER outperforms strong captioning, dual-encoder, visual matching and hierarchical classification baselines, affirming its advantage in tackling web-scale visual entity recognition. Our superior performance compared to dual encoders aligns with previous works observing that CLIP struggles for fine-grained recognition [12, 15]. Due to query image and entity name similarities being captured only through a vector dot product, potentially fine-grained interactions are missed. Also, GER offers significant advan-

Method	Entity-based pretraining			+ finetuning on seen		
	HM	seen	unseen	HM	seen	unseen
<i>Dual encoders</i>	9.2	8.9	9.4	16.3	24.3	12.3
<i>Visual matching</i>	16.2	15.5	17.1	16.4	15.7	17.2
<i>Captioning</i>	13.2	13.1	13.3	16.8	25.9	12.5
<i>Hierarchic. classif.</i>	14.7	14.8	14.6	21.8	29.6	17.2
<i>Generative entity recognition</i>						
GER-ATOMIC	15.9	15.3	16.7	20.1	26.2	16.3
GER-CAPTION	14.3	16.5	12.6	20.7	26.8	16.9
GER-HKC	15.8	15.5	16.0	21.0	25.2	17.9
GER-ALD	17.7	18.3	17.2	22.7	31.5	17.7

Table 2. **Baseline comparisons.** All baselines use exactly the *same pretraining dataset* Entity-WebLI (55M) and architectures of comparable number of parameters ($\sim 400M$). All numbers are obtained with finetuning on seen split after entity-based pretraining. We report the Harmonic Mean of top-1 accuracy on OVEN test.

tages over dual encoders: its computational complexity is not a function of entity set size and it does not require to store entity dense embeddings.

Different GER variants. In Tab. 2, we compare different variants of GER: one variant using unstructured codes (GER-ATOMIC) and three variants using semantically-structured codes: GER-CAPTION, GER-HKC and GER-ALD. We observe that GER-ALD is the best performing variant, both after entity-based pretraining and after finetuning on the OVEN seen entities. Compared to GER-CAPTION, GER-ALD use codes that are more discriminative and compact, which improves the performance particularly for entities with long names (see Sec. 4.4.2). Compared to GER-ATOMIC, GER-ALD codes yield a semantic structure which is crucial for million-scale label-space as shown in Sec. 4.4.1. GER-HKC model also gets strong performance but relies on an off-the-shelf semantic text encoder which makes the approach more complex and costly compared to GER-ALD. GER-HKC is a first step towards *learning* codes and we hope future works will propose original and better code creation strategies [41].

4.4. Analysis and ablation study

In this section, unless specified otherwise, we report the accuracy on the OVEN validation set [12] evaluated after pretraining on Entity-WebLI (27M), *i.e.* no OVEN finetuning.

4.4.1 Semantic versus atomic codes

In Fig. 3 (and Appendix Tab. 6), we report the relative improvement of semantically-structured codes (GER-ALD) compared to unstructured codes (GER-ATOMIC). We vary pretraining data size, model capacity and label-space size. A relative improvement of 100% means that the performance of GER-ALD doubles compared to GER-ATOMIC.

Limited pretraining data. In Fig. 3 (left), we see that semantic codes outperform atomic codes when the amount of




Figure 3. **Semantic vs atomic codes.** We report the relative improvement in % of GER-ALD compared to GER-ATOMIC in 3 scenarios: (i) limited pretraining data, (ii) limited model capacity and (iii) massive-scale label-space. Plots share a common experiment shown by ■ which uses a pretraining dataset size of 27M, Large model and 6M entity set. The setting reported in Tab. 2 is ★.

data available for pretraining diminishes. In fact, the results reported in Tab. 2 corresponds to the most favorable scenario for GER-ATOMIC with 55M pretraining datapoints (represented by ★ in Fig. 3). The relative improvement in this case is still of 14% while it grows to more than 1000% when the amount of data is reduced by 5 \times .

Limited model capacity. In Fig. 3 (middle), we see that the model struggles to learn unstructured codes when its capacity is reduced. When considering the small version of our model (114M parameters), the performance with atomic codes is very poor: 0.7 top-1 accuracy.

Web-scale label-space. In Fig. 3 (right), we vary the number of entities for pretraining. The “positive” entities (see Sec. 4.1) are always included in the pretraining set and the amount of “negative” entities is increased, effectively acting as distractors. First, we see in Fig. 3 (right) that for relatively small-scale label-space ($\leq 100k$), the benefit of having semantic codes versus atomic is small. In this regime we find that the model can memorize all the entities without the need for semantic structure between them. This aligns with the findings of DSI [42]. We evaluate GER further in small label-spaces in Sec. 4.5. However, we see that in million-scale label-space regime, semantic structure becomes important and significantly improves the performance compared to atomic codes: +26% relative improvement.

Overall, we find that GER-ATOMIC fail to learn unstructured codes when the amount of pretraining data or architecture capacity are reduced, or when the label-space increases to million-scale. Unlike GER-ATOMIC, GER-ALD succeed in these scenarios thanks to the semantic structure easing the learning. Next, we analyze how GER-ALD improves over another type of semantic codes: GER-CAPTION codes.

4.4.2 ALD versus captioning codes

We analyze why unambiguous, language-based and discriminative codes (GER-ALD) are more effective for entity recognition than directly decoding the entity name (GER-CAPTION). In Fig. 5 (left), we report the performance of GER-ALD and GER-CAPTION when varying the length L of




Figure 4. Accuracy per entity name length for GER-ALD versus GER-CAPTION codes. (left): Accuracy averaged per entity name length. (right): Qualitative examples of predictions for long entity names. Code tokens are symbolized between brackets.




Figure 5. ALD versus captioning codes. (left): Effect of different code lengths for GER-ALD and GER-CAPTION codes. (right): Cumulative distribution function (CDF) of (in green) the position of the least frequent token in the tokenized entity name and of (in pink) the length of tokenized entity name.

the codes. Fixing a code length L to a caption corresponds to keeping only the first L^{th} tokens of the entity name. In Fig. 5 (right), we report the cumulative distribution functions (CDF) of (i) the position within the tokenized entity name of the least frequent token among the entire corpus (as described in Sec. 3.2) and (ii) the total number of tokens in the tokenized entity name (L_e in the notations of Sec. 3.2)).

Discriminative tokens versus number of tokens. We observe in Fig. 5 (left) that the performance of GER-CAPTION increases drastically from $L = 2$ to $L = 4$. At the same time, we see in Fig. 5 (right) that for $L = 4$, less than half of the entity names are considered in full while more than 80% of the GER-CAPTION codes contain the least frequent token of the entire tokenized name. This hints that what is important for language-based codes is not to describe the full entity name but to include its most discriminative part. We also observe that the performance of captioning increases only moderately from $L = 4$ to $L = 8$ even though the number of entities considered in full increases drastically from 46.6% to 100%. This confirms our intuition that decoding all the entity name tokens does not have a major impact on the performance as long as the most discriminative tokens are decoded. Overall, these observations motivate the ALD design of keeping only the most discriminative tokens, which is shown in Fig. 5 to lead to improved performance compared to decoding the full tokenized entity name.

Effect of code length for GER-ALD. We see in Fig. 5 (left) that the performance of GER-ALD is the best for $L = 4$. With smaller code lengths, we need to resort to random tokens a lot to achieve unique codes (see Sec. 3.2), which deters the

Selection strategy	HM	Tokens order	HM
Least frequent tokens	14.4	Least frequent first	14.4
Most frequent tokens	12.3	Syntax order	14.4
First tokens	12.0	Random order	13.0
Random tokens	11.3	Least frequent last	12.7

Table 3. Ablation study of GER-ALD codes. (left) Word tokens selection. (right) Tokens order. All variants use $L = 4$. Default is in top rows. Non language-based GER-ATOMIC gets 11.4 top-1.

performance. For example at $L = 2$, more than 10% of the entities use a random code token while this percentage decreases to 0.5% at $L = 4$. We also see that the performance of GER-ALD decreases for code length above $L = 4$, which hints that only the few most discriminative tokens are important while additional ones clutter the entity code. Interestingly we also observe in Fig. 5 (left) that when considering all the tokens, GER-ALD performance is slightly below that of GER-CAPTION. This might seem surprising since the same amount of information is present in both cases. However we find that when considering all the tokens, it is more difficult for the model to decode tokens ordered by frequencies than tokens ordered syntactically.

Entities with long entity names. In Fig. 4 (left), we report the accuracy per entity name length for both GER-ALD and GER-CAPTION finetuned models. We see that the longer the entity name, the more GER-ALD improves over captioning. Longer entities tend to have more noise with key information further into the code. We also show in Fig. 4 qualitative examples of entities with long entity names (more in Fig. 12 in Appendix). In the left example, we see that GER-ALD use the token combination [col][ob] to represent the semantic concept of colobus monkey species. The last token is used to efficiently differentiate between sub-species of colobus. This compact and discriminative way of encoding the entity allows GER-ALD to successfully predict this entity whereas GER-CAPTION fails to generate the entity tokenized name.

4.4.3 Creating codes with ALD

Least frequent tokens. In Tab. 3 (left), we validate our choice of selecting the least frequent tokens by evaluating 3 alternatives: random choice, most frequent tokens and first-appearing tokens in tokenized entity name. We see that these alternative strategies hurt the performance significantly. Qualitative examples in Appendix Fig. 11 show that

Dataset	Codes	HM
WebLI	WebLI caption	1.8
Entity-WebLI (55M)	WebLI caption	12.9 (+11.1)
Entity-WebLI (55M)	Entity name	14.8 (+1.9)
Entity-WebLI (55M)	ALD	17.5 (+2.7)



Figure 6. **Entity-based pretraining ablation.** (left): Validation OVEN accuracy. (right): Examples of original WebLI captions versus corresponding OVEN entity names.




Figure 7. **Pretraining.** We vary the size of the pretraining dataset by changing the amount of retrieved examples from WebLI for each OVEN entity (see Sec. 3.3).

the kept tokens are less semantic and discriminative compared to GER-ALD strategy of keeping the least frequent tokens. Note that all these variants are at least as good as GER-ATOMIC (11.4 top-1) which is not based on language at all.

Decoding order. In Tab. 3 (right), we vary the order of the first $L - 1$ tokens in GER-ALD codes. Instead of decoding tokens from least to most frequent, we evaluate most to least frequent, syntax order and random order. Note that the selected tokens are the same in all variants, only their order changes. We see that both “least frequent first” and “syntax” orders achieve the best of performance.

4.4.4 Entity-based pretraining

Entity-based pretraining. In Fig. 6, we analyze why our entity-based pretraining improves over the standard captioning pretraining of PaLI or GiT models. First, we see that our method of selecting WebLI data relevant to OVEN entities drastically improves the performance (+11.1 in Fig. 6 (left)). This is because, by design, we select image-text pairs from WebLI that have captions similar to OVEN entity names. Hence, this data is directly relevant for the OVEN entity recognition benchmark. Second, we see that replacing the original WebLI caption with its corresponding entity name from OVEN leads to superior performance (+1.9). We see in the qualitative examples of Fig. 6 (right) that original captions contain a lot of descriptive information not directly relevant to the entity. Lastly, we confirm that using GER-ALD codes is better (+2.7) than tokenized entity name.

Dataset size. In Fig. 7, we evaluate the effect of the pretraining dataset size for GER models. We control the dataset size by varying the amount of retrieved examples from WebLI for each of the OVEN entities (see Sec. 3.3). We see in Fig. 7 that GER-ALD, GER-CAPTION and GER-ATOMIC benefit greatly from more data and do not seem to have reached saturation yet. As analyzed in Sec. 4.4.1, GER-ATOMIC fails when the amount of pretraining data decreases.

Method	ImageNet-LT	WebVision
Classif. MLP	74.3	80.9
GER-ATOMIC $L = 2$	80.8	84.7
GER-ALD $L = 2$	80.9	84.8
GER-ATOMIC $L = 1$ (\sim Classif. MAP)	81.0	84.8

Previously published numbers

NCR [13]	–	76.8
CurrNet [10]	–	79.3
PEL [40]	78.3	–
MAM [14]†	82.3	83.6

Table 4. **Evaluation of classification models and GER on small-scale label-spaces.** † indicates the use of additional data.

4.5. Link with classification

A typical way of tackling visual entity recognition is by training a classifier into the number of entities [35]. This is not a viable solution for web-scale problems such as OVEN where a single fully-connected layer for a 6M classes has an enormous parameter count of 4.6B. In this section, we evaluate GER in cases where learning a classification model is a feasible choice (smaller number of classes). Classification can be cast in our GER framework simply by setting $L = 1$ and $V = |\mathcal{E}| = \text{number of classes}$ (see Sec. 3.1), making it a special case of atomic codes with $L = 1$. Since the decoder decodes a single token, it is equivalent to a multi-layer Multihead Attention Pooling (MAP) head [21, 48]. In Tab. 4, we consider two challenging classification datasets: long-tailed ImageNet-LT [24] and noisy Webvision [22]. We evaluate GER- $\{\text{ALD, ATOMIC}\}$ and a classification baseline using multi-layer perceptron (MLP) on averaged-pooled patch tokens. Implementation details are in Sec 6.3 in Appendix.

We see in Tab. 4 that using GER-ATOMIC instead of standard MLP improves significantly the performance of the classification model (74.3 versus 81.0 for ImageNet-LT). We also observe that GER-ATOMIC and GER-ALD have comparable performance in this relatively small label-space regime (1k classes). As a matter of fact, this achieves state-of-the-art accuracy for both datasets (when no additional external data is used). This shows that GER framework not only excels for large-scale scenarios, but also works well in datasets with smaller number of visual entities, making GER a general framework for visual entity recognition.

5. Conclusion

In this work, we propose a novel generative framework for web-scale visual entity recognition. We represent each entity by a compact, discriminative and semantic code that a generative auto-regressive model learns to decode. In future work, we will explore ways of creating better entity codes by leveraging additional information: either from the Wikipedia page such as the description of the entity and its attached image or also by using external tools.

References

- [1] Rahul Agrawal, Archit Gupta, Yashoteja Prabhu, and Manik Varma. Multi-label learning with millions of labels: Recommending advertiser bid phrases for web pages. In *Proceedings of the 22nd international conference on World Wide Web*, pages 13–24, 2013. [2](#)
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. [1](#)
- [3] Samy Bengio, Krzysztof Dembczynski, Thorsten Joachims, Marius Kloft, and Manik Varma. Extreme Classification (Dagstuhl Seminar 18291). *Dagstuhl Reports*, 2019. [2](#)
- [4] Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101—mining discriminative components with random forests. In *ECCV*, 2014. [2](#)
- [5] Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. *ICLR*, 2023. [1](#), [2](#), [4](#), [5](#), [11](#), [12](#)
- [6] Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. Autoregressive entity retrieval. *arXiv preprint arXiv:2010.00904*, 2020. [1](#), [2](#), [3](#), [5](#)
- [7] Mostafa Dehghani, Alexey Gritsenko, Anurag Arnab, Matthias Minderer, and Yi Tay. Scenic: A JAX library for computer vision research and beyond. *arXiv preprint arXiv:2110.11403*, 2021. [5](#)
- [8] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 88, 2010. [2](#)
- [9] Li Fei-Fei, Rob Fergus, and Pietro Perona. Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. In *CVPR*, 2004. [2](#)
- [10] Sheng Guo, Weilin Huang, Haozhi Zhang, Chenfan Zhuang, Dengke Dong, Matthew R Scott, and Dinglong Huang. CurriculumNet: Weakly supervised learning from large-scale web images. In *ECCV*, pages 135–150, 2018. [8](#)
- [11] Michael Gutmann and Aapo Hyvärinen. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics. JMLR Workshop and Conference Proceedings*, 2010. [1](#), [5](#)
- [12] Hexiang Hu, Yi Luan, Yang Chen, Urvashi Khandelwal, Mandar Joshi, Kenton Lee, Kristina Toutanova, and Ming-Wei Chang. Open-domain visual entity recognition: Towards recognizing millions of wikipedia entities. *ICCV*, 2023. [1](#), [2](#), [3](#), [5](#), [6](#)
- [13] Ahmet Iscen, Jack Valmadre, Anurag Arnab, and Cordelia Schmid. Learning with neighbor consistency for noisy labels. In *CVPR*, 2022. [8](#)
- [14] Ahmet Iscen, Alireza Fathi, and Cordelia Schmid. Improving image recognition by retrieving from web-scale image-text data. *CVPR*, 2023. [8](#)
- [15] Ahmet Iscen, Mathilde Caron, Alireza Fathi, and Cordelia Schmid. Retrieval-enhanced contrastive vision-text models. *ICLR*, 2024. [1](#), [2](#), [4](#), [5](#)
- [16] 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. [5](#)
- [17] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Li Fei-Fei. Novel dataset for fine-grained image categorization. In *First Workshop on Fine-Grained Visual Categorization, CVPR*, 2011. [2](#)
- [18] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV*, 2013. [2](#)
- [19] Taku Kudo. Subword regularization: Improving neural network translation models with multiple subword candidates. *arXiv preprint arXiv:1804.10959*, 2018. [3](#)
- [20] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. *arXiv preprint arXiv:1808.06226*, 2018. [2](#), [3](#), [14](#)
- [21] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, 2019. [8](#)
- [22] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017. [2](#), [8](#), [11](#)
- [23] Haotian Liu, Kilho Son, Jianwei Yang, Ce Liu, Jianfeng Gao, Yong Jae Lee, and Chunyuan Li. Learning customized visual models with retrieval-augmented knowledge. In *CVPR*, 2023. [2](#), [4](#)
- [24] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X. Yu. Large-scale long-tailed recognition in an open world. In *CVPR*, 2019. [2](#), [8](#), [11](#)
- [25] Sanket Vaibhav Mehta, Jai Gupta, Yi Tay, Mostafa Dehghani, Vinh Q Tran, Jinfeng Rao, Marc Najork, Emma Strubell, and Donald Metzler. Dsi++: Updating transformer memory with new documents. *arXiv preprint arXiv:2212.09744*, 2022. [1](#), [2](#)
- [26] Anshul Mittal, Kunal Dahiya, Shreya Malani, Janani Ramaswamy, Seba Kuruvilla, Jitendra Ajmera, Keng-hao Chang, Sumeet Agarwal, Purushottam Kar, and Manik Varma. Multi-modal extreme classification. In *CVPR*, 2022. [2](#)
- [27] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in neural information processing systems*, 32, 2019. [4](#)
- [28] Jianmo Ni, Gustavo Hernández Ábrego, Noah Constant, Ji Ma, Keith B Hall, Daniel Cer, and Yinfei Yang. Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models. *arXiv preprint arXiv:2108.08877*, 2021. [4](#), [12](#)
- [29] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [1](#), [5](#)

- [30] OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. [1](#)
- [31] Ronak Pradeep, Kai Hui, Jai Gupta, Adam D Lelkes, Honglei Zhuang, Jimmy Lin, Donald Metzler, and Vinh Q Tran. How does generative retrieval scale to millions of passages? *arXiv preprint arXiv:2305.11841*, 2023. [1](#), [2](#)
- [32] 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](#), [4](#), [5](#)
- [33] Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan H Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Q Tran, Jonah Samost, et al. Recommender systems with generative retrieval. *arXiv preprint arXiv:2305.05065*, 2023. [1](#), [2](#)
- [34] Stephen Robertson, Hugo Zaragoza, et al. The probabilistic relevance framework: Bm25 and beyond. *Foundations and Trends® in Information Retrieval*, 2009. [2](#)
- [35] 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. *International journal of computer vision*, 2015. [2](#), [8](#)
- [36] Florian Schroff, Dmitry Kalenichenko, and James Philbin. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015. [5](#)
- [37] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [4](#)
- [38] 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. *arXiv preprint arXiv:2210.08402*, 2022. [4](#)
- [39] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*, 2015. [3](#)
- [40] Jiang-Xin Shi, Tong Wei, Zhi Zhou, Xin-Yan Han, Jie-Jing Shao, and Yu-Feng Li. Parameter-efficient long-tailed recognition. *arXiv preprint arXiv:2309.10019*, 2023. [8](#)
- [41] Weiwei Sun, Lingyong Yan, Zheng Chen, Shuaiqiang Wang, Haichao Zhu, Pengjie Ren, Zhumin Chen, Dawei Yin, Maarten Rijke, and Zhaochun Ren. Learning to tokenize for generative retrieval. *NeurIPS*, 2023. [1](#), [2](#), [6](#)
- [42] Yi Tay, Vinh Tran, Mostafa Dehghani, Jianmo Ni, Dara Bahri, Harsh Mehta, Zhen Qin, Kai Hui, Zhe Zhao, Jai Gupta, et al. Transformer memory as a differentiable search index. *Advances in Neural Information Processing Systems*, 2022. [1](#), [2](#), [5](#), [6](#)
- [43] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *CVPR*, 2018. [2](#)
- [44] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011. [2](#)
- [45] 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. *arXiv preprint arXiv:2205.14100*, 2022. [3](#), [4](#), [5](#), [11](#), [12](#)
- [46] Yujing Wang, Yingyan Hou, Haonan Wang, Ziming Miao, Shibin Wu, Qi Chen, Yuqing Xia, Chengmin Chi, Guoshuai Zhao, Zheng Liu, et al. A neural corpus indexer for document retrieval. *NeurIPS*, 2022. [2](#), [4](#)
- [47] Tobias Weyand, Andre Araujo, Bingyi Cao, and Jack Sim. Google landmarks dataset v2-a large-scale benchmark for instance-level recognition and retrieval. In *CVPR*, 2020. [2](#)
- [48] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *CVPR*, 2022. [8](#)
- [49] Yidan Zhang, Ting Zhang, Dong Chen, Yujing Wang, Qi Chen, Xing Xie, Hao Sun, Weiwei Deng, Qi Zhang, Fan Yang, et al. Irgen: Generative modeling for image retrieval. *arXiv preprint arXiv:2303.10126*, 2023. [2](#)
- [50] Zheng Zhu, Guan Huang, Jiankang Deng, Yun Ye, Junjie Huang, Xinze Chen, Jiagang Zhu, Tian Yang, Dalong Du, Jiwen Lu, et al. Webface260m: A benchmark for million-scale deep face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [2](#)