

The Curious Layperson: Fine-Grained Image Recognition without Expert Labels

Subhabrata Choudhury
subha@robots.ox.ac.uk

Iro Laina
iro@robots.ox.ac.uk

Christian Rupprecht
chrshr@robots.ox.ac.uk

Andrea Vedaldi
vedaldi@robots.ox.ac.uk

Visual Geometry Group
University of Oxford
Oxford, UK

Abstract

Most of us are not experts in specific fields, such as ornithology. Nonetheless, we do have general image and language understanding capabilities that we use to match what we see to expert resources. This allows us to expand our knowledge and perform novel tasks without ad-hoc external supervision. On the contrary, machines have a much harder time consulting expert-curated knowledge bases unless trained specifically with that knowledge in mind. Thus, in this paper we consider a new problem: fine-grained image recognition without expert annotations, which we address by leveraging the vast knowledge available in web encyclopedias. First, we learn a model to describe the visual appearance of objects using non-expert image descriptions. We then train a fine-grained textual similarity model that matches image descriptions with documents on a sentence-level basis. We evaluate the method on two datasets and compare with several strong baselines and the state of the art in cross-modal retrieval. Code is available at: <https://github.com/subhc/clever>.

1 Introduction

Deep learning and the availability of large-scale labelled datasets have led to remarkable advances in image recognition tasks, including fine-grained recognition [21, 36, 57]. The problem of fine-grained image recognition amounts to identifying subordinate-level categories, such as different species of birds, dogs or plants. Thus, the supervised learning regime in this case requires annotations provided by domain *experts* or citizen scientists [52].

While most people, unless professionally trained or enthusiasts, do not have knowledge in such specific domains, they are generally capable of consulting existing expert resources such as books or online encyclopedias, *e.g.* Wikipedia. As an example, let us consider bird identification. Amateur bird watchers typically rely on field guides to identify observed species. As a general instruction, one has to answer the question “what is most noticeable about this bird?” before skimming through the guide to find the best match to their observation.

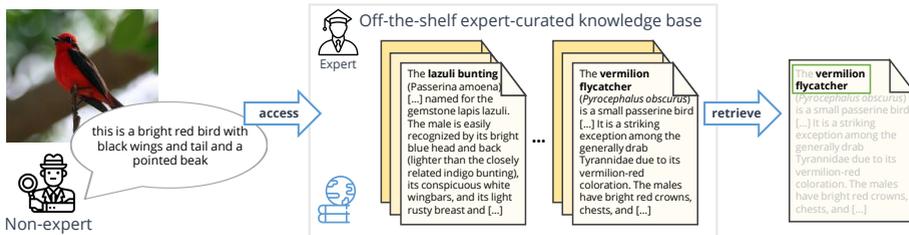


Figure 1: Fine-Grained Image Recognition without Expert Labels. We propose a novel task that enables fine-grained classification without using expert class information (e.g. bird species) during training. We frame the problem as document retrieval from general image descriptions by leveraging existing textual knowledge bases, such as Wikipedia.

The answer to this question is typically a detailed description of the bird’s shape, size, plumage colors and patterns. Indeed, in Fig. 1, the non-expert observer might not be able to directly identify a bird as a “Vermillion Flycatcher”, but they *can* simply describe the appearance of the bird: “*this is a bright red bird with black wings and tail and a pointed beak*”. This description can be matched to an expert corpus to obtain the species and other expert-level information.

On the other hand, machines have a much harder time consulting off-the-shelf expert-curated knowledge bases. In particular, most algorithmic solutions are designed to address a *specific* task with datasets constructed *ad-hoc* to serve precisely this purpose. Our goal, instead, is to investigate whether it is possible to re-purpose general image and text understanding capabilities to allow machines to consult already existing *textual* knowledge bases to address a new task, such as recognizing a bird.

We introduce a novel task inspired by the way a layperson would tackle fine-grained recognition from visual input; we name this **CLEVER**, *i.e.* Curious Layperson-to-Expert Visual Entity Recognition. Given an image of a subordinate-level object category, the task is to retrieve the relevant document from a large, expertly-curated text corpus; to this end, we only allow non-expert supervision for learning to describe the image. We assume that: (1) the corpus dedicates a separate entry to each category, as is, for example, the case in encyclopedia entries for bird or plant species, etc., (2) there exist no paired data of images and documents or expert labels during training, and (3) to model a layperson’s capabilities, we have access to general image and text understanding tools that do not use expert knowledge, such as image descriptions or language models.

Given this definition, the task classifies as weakly-supervised in the taxonomy of learning problems. We note that there are fundamental differences to related topics, such as image-to-text retrieval and unsupervised image classification. Despite a significant amount of prior work in image-to-text or text-to-image retrieval [20, 22, 41, 53, 72], the general assumption is that images and corresponding documents are paired for training a model. In contrast to unsupervised image classification, the difference is that here we are interested in *semantically* labelling images using a secondary modality, instead of grouping similar images [6, 8, 61].

To the best of our knowledge, we are the first to tackle the task of fine-grained image recognition without expert supervision. Since the target corpus is not required during training, the search domain is easily extendable to any number of categories/species—an ideal use case when retrieving documents from dynamic knowledge bases, such as Wikipedia. We provide extensive evaluation of our method and also compare to approaches in cross-modal retrieval,

despite using significantly reduced supervision.

2 Related Work

In this paper, we address a novel problem (CLEVER). Next we describe in detail how it differs from related problems in the computer vision and natural language processing literature and summarise the differences with respect to how class information is used in Table 1.

Fine-Grained Recognition. The goal of fine-grained visual recognition (FGVR) is categorising objects at subordinate level, such as species of animals or plants [29, 67, 62, 63, 62]. Large-scale annotated datasets require domain experts and are thus difficult to collect. FGVR is more challenging than coarse-level image classification as it involves categories with fewer discriminative cues and fewer labeled samples. To address this problem, supervised methods exploit side information such as part annotations [71], attributes [65], natural language descriptions [19], noisy web data [18, 28, 69] or humans in the loop [7, 9, 10]. Attempts to reduce supervision in FGVR are mostly targeted towards eliminating auxiliary labels, *e.g.* part annotations [71, 24, 49, 73]. In contrast, our goal is fine-grained recognition without access to *categorical* labels during training. Our approach only relies on side information (captions) provided by laymen and is thus unsupervised from the perspective of “expert knowledge”.

Zero/Few Shot Learning. Zero-shot learning (ZSL) is the task of learning a classifier for unseen classes [65]. A classifier is generated from a description of an object in a secondary modality, mapping semantic representations to class space in order to recognize said object in images [60]. Various modalities have been used as auxiliary information: word embeddings [16, 62], hierarchical embeddings [26], attributes [8, 24] or Wikipedia articles [12, 13, 44, 72]. Most recent work uses generative models conditioned on class descriptions to synthesize training examples for unseen categories [15, 27, 62, 66, 66, 67]. The multi-modal and often fine-grained nature of the standard and generalised (G)ZSL task renders it related to our problem. However, different from the (G)ZSL settings our method uses neither class supervision during training nor image-document pairs as in [12, 13, 44, 72].

Cross-Modal and Information Retrieval. While information retrieval deals with extracting information from document collections [65], cross-modal retrieval aims at retrieving relevant information across various modalities, *e.g.* image-to-text or vice versa. One of the core problems in information retrieval is ranking documents given some query, with a classical example being Okapi BM25 [48]. With the advent of transformers [64] and BERT [10], state-of-the-art document retrieval is achieved in two-steps; an initial ranking based on keywords followed by computationally intensive BERT-based re-ranking [64, 38, 39, 70]. In cross-modal retrieval, the common approach is to learn a shared representation space for multiple modalities [9, 20, 22, 40, 41, 42, 58, 62, 72]. In addition to paired data in various domains, some methods also exploit auxiliary semantic labels; for example, the Wikipedia benchmark [43] provides broad category labels such as *history*, *music*, *sport*, etc.

We depart substantially from the typical assumptions made in this area. Notably, with the exception of [20, 60], this setting has not been explored in fine-grained domains, but generally

Task	Class Information	
	Train	Test
FGVR	K	K
ZSL	K	U
GZSL	K	K + U
CLEVER	U	U

Table 1: Overview of related topics (K: known, U: unknown).

targets higher-level content association between images and documents. Furthermore, one major difference between our approach and cross-modal retrieval, including [20, 60], is that we do not assume paired data between the input domain (images) and the target domain (documents). We address the lack of such pairs using an intermediary modality (captions) that allows us to perform retrieval directly in the text domain.

Natural Language Inference (NLI) and Semantic Textual Similarity (STS). Also related to our work, in natural language processing, the goal of the NLI task is to recognize textual entailment, *i.e.* given a pair of sentences (premise and hypothesis), the goal is to label the hypothesis as entailment (true), contradiction (false) or neutral (undetermined) with respect to the premise [6, 63]. STS measures the degree of semantic similarity between two sentences [11, 2]. Both tasks play an important role in semantic search and information retrieval and are currently dominated by the transformer architecture [10, 31, 47, 54]. Inspired by these tasks, we propose a sentence similarity regime that is domain-specific, paying attention to fine-grained semantics.

3 Method

We introduce the problem of layperson-to-expert visual entity recognition (CLEVER), which we address via image-based document retrieval. Formally, we are given a set of images $x_i \in \mathcal{I}$ to be labelled given a corpus of expert documents $D_j \in \mathcal{D}$, where each document corresponds to a fine-grained image category and there exist $K = |\mathcal{D}|$ categories in total. As a concrete example, \mathcal{I} can be a set of images of various bird species and \mathcal{D} a bird identification corpus constructed from specialized websites (with one article per species). Crucially, the pairing of x_i and D_j is not known, *i.e.* no expert task supervision is available during training. Therefore, the mapping from images to documents cannot be learned directly but can be discovered through the use of non-expert image descriptions C_i for image x_i .

Our method consists of three distinct parts. First, we learn, using “layperson’s supervision”, an image captioning model that uses simple color, shape and part descriptions. Second, we train a model for Fine-Grained Sentence Matching (FGSM). The FGSM model takes as input a pair of sentences and predicts whether they are descriptions of the same object. Finally, we use the FGSM to score the documents in the expert corpus via voting. As there is one document per class, the species corresponding to the highest-scoring document is returned as the final class prediction for the image. The overall inference process is illustrated in Fig. 2.

3.1 Fine-grained Sentence Matching

The overall goal of our method is to match images to expert documents — however, in absence of paired training data, learning a cross-domain mapping is not possible. On the other hand, describing an image is an easy task for most humans, as it usually does not require domain knowledge. It is therefore possible to leverage image descriptions as an intermediary for learning to map images to an expert corpus.

To that end, the core component of our approach is the FGSM model $f(c_1, c_2) \in \mathbb{R}$ that scores the visual similarity of two descriptions c_1 and c_2 . We propose to train f in a manner similar to the textual entailment (NLI) task in natural language processing. The difference to NLI is that the information that needs to be extracted here is fine-grained and domain-specific *e.g.* “*a bird with blue wings*” vs. “*this is a uniformly yellow bird*”. Since we do not have annotated sentence pairs for this task, we have to create them synthetically. Instead of the

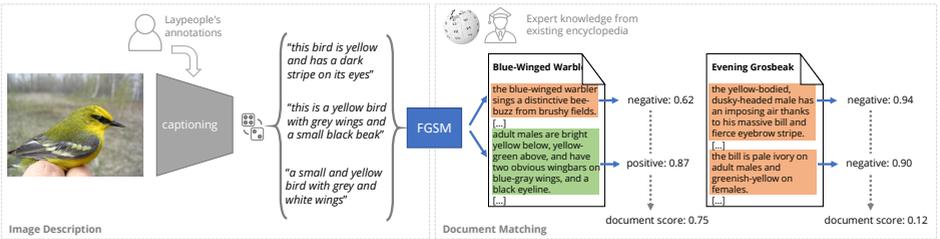


Figure 2: **Overview.** We train a model for fine-grained sentence matching (FGSM) using layperson’s annotations, *i.e.* class-agnostic image descriptions. At test time, we score documents from a relevant corpus and use the top-ranked document to label the image.

terms entailment and contradiction, here we use positive and negative to emphasize that the goal is to find matches (or mismatches) between image descriptions.

We propose to model f as a sentence encoder, performing the semantic comparison of c_1, c_2 in embedding space. Despite their widespread success in downstream tasks, most transformer-based language models are notoriously bad at producing semantically meaningful sentence embeddings [80, 47]. We thus follow [47] in learning an appropriate textual similarity model with a Siamese architecture built on a pre-trained language transformer. This also allows us to leverage the power of large language models while maintaining efficiency by computing an embedding for each input independently and only compare embeddings as a last step. To this end, we compute a similarity score for c_1 and c_2 as $f(c_1, c_2) = h([\phi_1; \phi_2; |\phi_1 - \phi_2|])$, where $[\cdot]$ denotes concatenation, and h and ϕ are lightweight MLPs operating on the average-pooled output of a large language model $T(\cdot)$ with the shorthand notation $\phi_1 = \phi(T(c_1))$.

Training. One requirement is that the FGSM model should be able to identify *fine-grained* similarities between pairs of sentences. This is in contrast to the standard STS and NLI tasks in natural language understanding which determine the relationship (or degree of similarity) of a sentence pair on a *coarser* semantic level. Since our end-goal is visual recognition, we instead train the model to emphasize visual cues and nuanced appearance differences.

Let C_i be the set of human-annotated descriptions for a given image x_i . Positive training pairs are generated by exploiting the fact that, commonly, each image has been described by multiple annotators; for example in CUB-200 [57] there are $|C_i| = 10$ captions per image. Thus, each pair (from $C_i \times C_i$) of descriptions of the same image can be used as a positive pair. The negative counterparts are then sampled from the complement $\bar{C}_i = \bigcup_{l \neq i} C_l$, *i.e.* among the available descriptions for all other images in the dataset. We construct this dataset with an equal amount of samples for both classes and train f with a binary cross entropy loss.

Inference. During inference the sentence embeddings ϕ for each sentence in each document can be precomputed and only h needs to be evaluated dynamically given an image and its corresponding captions, as described in the next section. This greatly reduces the memory and time requirements.

3.2 Document Scoring

Although trained from image descriptions alone, the FGSM model can take any sentence as input and, at test time, we use the trained model to score sentences from an expert corpus

against image descriptions. Specifically, we assign a score $z_{ij} \in \mathbb{R}$ to each expert document D_j given a set of descriptions for the i -th image:

$$z_{ij} = \frac{1}{|\mathcal{C}_i \times D_j|} \sum_{(c,s) \in \mathcal{C}_i \times D_j} f(c, s), \quad (1)$$

Since there are several descriptions in \mathcal{C}_i and sentences in D_j , we compute the final document score as an average of individual predictions (scores) of all pairs of descriptions and sentences. Aggregating scores across the whole corpus \mathcal{D} , we can then compute the probability $p(D_j | x_i) = \frac{e^{-z_{ij}}}{\sum_k e^{-z_{ik}}}$ of a document $D_j \in \mathcal{D}$ given image x_i and assign the document (and consequently class) with the highest probability to the image.

3.3 Bridging the Domain Gap

While training the FGSM model, we have so far only used laypersons’ descriptions, disregarding the expert corpus. However, we can expect the documents to contain significantly more information than visual descriptions. In the case of bird species, encyclopedia entries usually also describe behavior, migration, conservation status, etc. In this section, we thus employ two mechanisms to bridge the gap between the image descriptions and the documents.

Neutral Sentences. We introduce a third, *neutral* class to the classification problem, designed to capture sentences that do not provide relevant (visual) information. We generate neutral training examples by pairing an image description with sentences from the documents (or other descriptions) that do not have any *nouns* in common. Instead of binary cross entropy, we train the three-class model (positive/neutral/negative) with softmax cross entropy.

Score Distribution Prior. Despite the absence of paired training data, we can still impose priors on the document scoring. To this end, we consider the probability distribution $p(\mathcal{D} | x)$ over the entire corpus \mathcal{D} given an image x in a training batch \mathcal{B} . We can then derive a regularizer $R(\mathcal{B})$ that operates at batch-level:

$$R(\mathcal{B}) = \sum_{x \in \mathcal{B}} \left(-\langle p(\mathcal{D} | x), p(\mathcal{D} | x) \rangle + \sum_{x' \in \mathcal{B}, x} \langle p(\mathcal{D} | x), p(\mathcal{D} | x') \rangle \right) \quad (2)$$

where $\langle \cdot, \cdot \rangle$ denotes the inner product of two vectors. The intuition of the two terms of the regularizer is as follows. $\langle p(\mathcal{D} | x), p(\mathcal{D} | x) \rangle$ is maximal when the distribution assigns all mass to a single document. Since the score z_{ij} is averaged over all captions of one image, this additionally has the side effect of encouraging all captions of one image to vote for the same document. The second term of $R(\mathcal{B})$ then encourages the distributions of two different images to be orthogonal, favoring the assignment of images uniformly across all documents.

Since $R(\mathcal{B})$ requires evaluation over the whole document corpus for every image, we first pre-train f , including the large transformer model T , (c.f. Section 3.1). After convergence, we extract sentence features for all documents and image descriptions and train only the MLPs ϕ and h with $\mathcal{L} + \lambda R$, where λ balances the 3-class cross entropy loss \mathcal{L} and the regularizer.

4 Experiments

We validate our method empirically for bird and plant identification. To the best of our knowledge, we are the first to consider this task, thus in absence of state-of-the-art methods, we ablate the different components of our model and compare to several strong baselines.

Method	CUB-200			FLO		
	top-1↑	top-5↑	MR↓	top-1↑	top-5↑	MR↓
random guess	0.5	2.5	100.0	0.9	4.9	51.0
SRoBERTa-STsb [47] (no-ft)	1.3	6.4	73.4	1.1	7.7	45.2
SRoBERTa-NLI [30] (no-ft)	1.9	5.3	81.3	0.9	5.7	48.2
Okapi BM25 [48]	1.0	7.5	78.2	1.6	8.0	43.9
TF-IDF [25]	2.2	9.7	72.1	1.4	5.0	45.2
RoBERTa [50]	4.3	16.6	44.6	1.1	9.6	42.6
ours	7.9	28.6	31.9	6.2	14.2	39.7

Table 2: **Comparison to baselines.** We report the retrieval performance of our method on CUB-200 and Oxford-102 Flowers (FLO) and compare to various strong baselines.

4.1 Datasets and Experimental Setup

Datasets. We evaluate our method on Caltech-UCSD Birds-200-2011 (CUB-200) [57] and the Oxford-102 Flowers (FLO) dataset [36]. For both datasets, Reed et al. [46] have collected several visual descriptions per image by crowd-sourcing to non-experts on Amazon Mechanical Turk (AMT). We further collect for each class a corresponding expert document from specialised websites, such as AllAboutBirds¹ (AAB) and Wikipedia.

Setup. We use the image-caption pairs to train two image captioning models: “Show, Attend and Tell” (SAT) [68] and AoANet [23]. Unless otherwise specified, we report the performance of our model based on their ensemble, *i.e.* combining predictions from both models. As the backbone T of our sentence transformer model, we use RoBERTa-large [50] fine-tuned on NLI and STS datasets using the setup of [47]. Please see the appendix for further implementation, architecture, dataset and training details.

We use three metrics to evaluate the performance on the benchmark datasets. We compute top-1 and top-5 per-class retrieval accuracy and report the overall average. Additionally, we compute the mean rank (MR) of the target document for each class. Here, retrieval accuracy is identical to classification accuracy, since there is only a single relevant article per category.

4.2 Baseline Comparisons

Since this work is the first to explore the mapping of images to expert documents without expert supervision, we compare our method to several strong baselines (Table 2).

Our FGSM performs text-based retrieval, we evaluate current text retrieval systems. **TF-IDF:** Term frequency-inverse document frequency (TF-IDF) is widely used for unsupervised document retrieval [25]. For each image, we use the predicted captions as queries and use the TF-IDF textual representation for document ranking instead of our model. We empirically found the cosine distance and n -grams with $n = 2, 3$ to perform best for TF-IDF. **BM25:** Similar to TF-IDF, BM25 [48] is another common measure for document ranking based on n -gram frequencies. We use the BM25 Okapi implementation from the python package `rank-bm25` with default settings. **RoBERTa:** One advantage of processing caption-sentence pairs with a Siamese architecture, such as SBERT/SRoBERTa [47], is the reduced complexity. Nonetheless, we have trained a transformer baseline for text classification, using

¹<https://allaboutbirds.com>

Method	top-1 \uparrow	top-5 \uparrow	MR \downarrow
user interaction	11.9	37.5	24.8
FGSM + cosine	4.5	17.8	35.5
FGSM w/ SAT	4.3	15.0	42.9
FGSM w/ AoANet	5.7	20.8	38.3
FGSM w/ ensemble	5.9	20.0	36.1
FGSM + $R(\mathcal{B})$ [2-cla]	7.4	24.6	29.9
FGSM + $R(\mathcal{B})$ [3-cla]	7.9	28.6	31.9

Table 3: **Ablation and user study.** On CUB-200 we evaluate scoring functions, captioning models and the regularizer $R(\mathcal{B})$.

the same backbone [53], concatenating each sentence pair with a SEP token and training as a binary classification problem. We apply this model to score documents, instead of FGSM, aggregating scores at sentence-level. **SROBERTa-NLI/STSB**: Finally, to evaluate the importance of learning *fine-grained* sentence similarities, we also measure the performance of the same model trained only on the NLI and STSB benchmarks [47], without further fine-tuning. Following [47] we rank documents based on the cosine similarity between the caption and sentence embeddings.

Our method outperforms all bag-of-words and learned baselines. Approaches such as TF-IDF and BM25 are very efficient, albeit less performant than learned models. Notably, the closest in performance to our model is the transformer baseline (RoBERTa), which comes at a large computational cost (347 sec vs. 0.55 sec for our model per image on CUB-200).

4.3 Ablation & User Interaction

We ablate the different components of our approach in Table 3. We first investigate the use of a different scoring mechanism, *i.e.* the cosine similarity between the embeddings of c and s as in [47]; we found this to perform worse (FGSM + cosine). We also study the influence of the captioning model on the final performance. We evaluate captions obtained by two methods, SAT [68] and AoANet [23], as well as their ensemble. The ensemble improves performance thanks to higher variability in the image descriptions. Next, we evaluate the performance of our model after the final training phase, with the proposed regularizer and the inclusion of neutral pairs (Section 3.3). $R(\mathcal{B})$ imposes prior knowledge about the expected class distribution over the dataset and thus stabilizes the training, resulting in improved performance ([2-cla]). Further, through the regularizer and neutral sentences ([3-cla]), FGSM is exposed to the target corpus during training, which helps reduce the domain shift during inference compared to training on image descriptions alone (FGSM w/ ensemble).

Finally, our method enables user interaction, *i.e.* allowing a user to directly enter own descriptions, replacing the automatic description model. In Table 3 we have simulated this by evaluating with ground-truth instead of predicted descriptions. Naturally, we find that human descriptions indeed perform better, though the performance gap is small. We attribute this gap to a much higher diversity in the human annotations. Current image captioning models still have diversity issues, which also explains why our ensemble variant improves the results.

Method	sup.	top-1 \uparrow	top-5 \uparrow	MR \downarrow
random guess	\times	2.0	10.0	25.0
ViLBERT [53]	\times	3.5	14.8	20.2
TF-IDF [45]	\times	7.2	28.6	18.9
CLIP [49]	\checkmark	10.0	32.9	14.0
DSCMR [72]	\checkmark	13.5	34.7	15.2
ours	\times	20.9	50.7	9.6

Table 4: **Comparison to cross-media retrieval.** We evaluate the performance of methods on the ZSL split of CUB-200. Our method performs favorably against existing approaches trained with more supervision.

Input	Predicted descriptions	Top 5 – retrieval results									
	<p>this bird has wings that are black and has a red belly this bird is red and black in color with a stubby red beak this bird has a red crown and a black eyering that is round this bird has wings that are black and has a red belly and black bill [...]</p>										
	<p>this bird has wings that are grey and has a white belly this bird is white with grey and has a long pointy beak this bird has wings that are white and has a black crown this bird has wings that are white and has a long orange beak [...]</p>										
	<p>this bird has wings that are black and has a red head this bird is black and red in color with a skinny black beak and black eye rings a small bird with a red breast and black and white feathers and a short beak this bird has wings that are black , white and has a red head [...]</p>										
	<p>this bird has wings that are brown and has a long neck this is a large grey bird with a large downward pointing beak this particular bird has a belly that is grey and white this is a large grey bird with a long neck and a large beak [...]</p>										
	<p>this is a small yellow bird with a grey head and a small pointy beak this little bird has a yellow belly and breast with a grey wing and white wingbar this bird has a yellow crown as well as a yellow belly this bird has wings that are black and has a yellow belly [...]</p>										

Figure 3: **Qualitative Results (CUB-200)**. We show examples of input images and their predicted captions, followed by the top-5 retrieved documents (classes). For illustration purposes, we show a random image for each document; the image is not used for matching.

4.4 Comparison with Cross-Modal Retrieval

Since the nature of the problem presented here is in fact cross-modal, we adapt a representative method, DSCMR [17], to our data to compare to the state of the art in cross-media retrieval. We note that such an approach requires image-document pairs as training samples, thus using more supervision than our method. Instead of using image descriptions as an intermediary for retrieval, DSCMR thus performs retrieval monolithically, mapping the modalities in a shared representation space. We argue that, although this is the go-to approach in broader category domains, it may be sub-optimal in the context of fine-grained categorization.

Since in our setting each category (species) is represented by a single article, in the scenario that a supervised model sees *all* available categories during training, the cross-modal retrieval problem degenerates to a classification task. Hence, for a meaningful comparison, we train both our model and DSCMR on the CUB-200 splits for ZSL [6] to evaluate on 50 *unseen* categories. We report the results in Table 4, including a TF-IDF baseline on the same split. Despite using no image-documents pairs for training, our method still performs significantly better.

Additionally, we compare to representative methods from the vision-and-language representation learning space. ViLBERT [8] is a multi-modal transformer model capable of learning joint representations of visual content and natural language. It is pre-trained on 3.3M image-caption pairs with two proxy tasks. We use their multi-modal alignment prediction mechanism to compute the alignment of the sentences in a document to a target image, similar to ViLBERT’s zero-shot experiments. The sentence scores are averaged to get the document alignment score and the document with the maximum score is chosen as the class. Finally, we compare to CLIP [45], that learns a multimodal embedding space from 400M image-text pairs. CLIP predicts image and sentence embeddings with separate encoders. For a target image we score each sentence using cosine similarity and average across the document for the final

score. CLIP’s training data is not public, but we find that there is a high possibility it does indeed contain expert labels as removing class names from documents hurts its performance.

4.5 Qualitative Results

In Fig. 3, we show qualitative retrieval results. The input image is shown on the left followed by the predicted descriptions. We then show the top-5 retrieved documents/classes together with an example image for the reader. Note that the example images are not used for matching, as the FGSM module operates on text only. We find that in most cases, even when the retrieved document does not match the ground truth class, the visual appearance is still similar. This is especially noticeable in families of birds for which discriminating among individual species is considered to be particularly difficult even for humans, *e.g.* warblers (last row).

5 Discussion

Like with any method that aims to reduce supervision, our method is not perfect. There are multiple avenues where our approach can be further optimized.

First, we observe that models trained for image captioning tend to produce short sentences that lack distinctiveness, focusing on the major features of the object rather than providing detailed fine-grained descriptions of the object’s unique aspects. We believe there is a scope for improvement if the captioning models could extensively describe each different part and attribute of the object. We have tried to mitigate this issue by using an ensemble of two popular captioning networks. However, using multiple models and sampling multiple descriptions may lead to redundancy. Devising image captioning models that produce diverse and distinct fine-grained image descriptions may provide improved performance on CLEVER task; there is an active area of research [69, 61] that is looking into this problem.

Second, the proposed approach to scoring a document given an image uses *all* the sentences in the document classifying them as positive, negative or neutral with respect to each input caption. Given that the information provided by an expert document might be noisy, *i.e.* not necessarily related to the *visual* domain, it is likely worthwhile to develop a filtering mechanism for relevancy, effectively using only a subset of the sentences for scoring.

Finally, in-domain regularization results in a significant performance boost (Table 3), which implies that the CLEVER task is susceptible to the domain gap between laypeople’s descriptions and the expert corpus. Language models such as BERT/RoBERTa partially address this problem already by learning general vocabulary, semantics and grammar during pre-training on large text corpora, enabling generalization to a new corpus without explicit training. However, further research in reducing this domain gap seems worthwhile.

6 Conclusion

We have shown that it is possible to address fine-grained image recognition without the use of expert training labels by leveraging existing knowledge bases, such as Wikipedia. This is the first work to tackle this challenging problem, with performance gains over the state of the art on cross-media retrieval, despite their training with image-document pairs. While humans can easily access and retrieve information from such knowledge bases, CLEVER remains a challenging learning problem that merits future research.

Acknowledgments

S. C. is supported by a scholarship sponsored by Facebook. I. L. is supported by the European Research Council (ERC) grant IDIU-638009 and EPSRC VisualAI EP/T028572/1. C. R. is supported by Innovate UK (project 71653) on behalf of UK Research and Innovation (UKRI) and ERC grant IDIU-638009. A. V. is supported by ERC grant IDIU-638009. We thank Andrew Brown for valuable discussions.

References

- [1] Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. SemEval-2012 task 6: A pilot on semantic textual similarity. In *SEM 2012*, pages 385–393, 7-8 June 2012.
- [2] Eneko Agirre, Daniel Cer, Mona Diab, Aitor Gonzalez-Agirre, and Weiwei Guo. *SEM 2013 shared task: Semantic textual similarity. In *SEM 2013*, pages 32–43, June 2013.
- [3] Zeynep Akata, Florent Perronnin, Zaid Harchaoui, and Cordelia Schmid. Label-embedding for image classification. *TPAMI*, 38(7):1425–1438, 2015.
- [4] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. Deep canonical correlation analysis. In *ICML*, pages 1247–1255. PMLR, 2013.
- [5] Yuki M. Asano, Christian Rupprecht, and Andrea Vedaldi. Self-labelling via simultaneous clustering and representation learning. In *ICLR, 2020*.
- [6] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. In *EMNLP*. Association for Computational Linguistics, 2015.
- [7] Steve Branson, Catherine Wah, Florian Schroff, Boris Babenko, Peter Welinder, Pietro Perona, and Serge Belongie. Visual recognition with humans in the loop. In *ECCV*, pages 438–451. Springer, 2010.
- [8] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. *NeurIPS*, 33, 2020.
- [9] Yin Cui, Feng Zhou, Yuanqing Lin, and Serge Belongie. Fine-grained categorization and dataset bootstrapping using deep metric learning with humans in the loop. In *CVPR*, pages 1153–1162, 2016.
- [10] Jia Deng, Jonathan Krause, Michael Stark, and Li Fei-Fei. Leveraging the wisdom of the crowd for fine-grained recognition. *TPAMI*, 38(4):666–676, 2015.
- [11] J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT, 2019*.
- [12] Mohamed Elhoseiny, Ahmed Elgammal, and Babak Saleh. Write a classifier: Predicting visual classifiers from unstructured text. *TPAMI*, 39(12):2539–2553, 2016.

- [13] Mohamed Elhoseiny, Yizhe Zhu, Han Zhang, and Ahmed Elgammal. Link the head to the "beak": Zero shot learning from noisy text description at part precision. In *CVPR*, pages 6288–6297. IEEE, 2017.
- [14] A. Farhadi, I. Endres, D. Hoiem, and D. Forsyth. Describing objects by their attributes. In *CVPR*, pages 1778–1785, 2009.
- [15] Rafael Felix, Vijay BG Kumar, Ian Reid, and Gustavo Carneiro. Multi-modal cycle-consistent generalized zero-shot learning. In *ECCV*, pages 21–37, 2018.
- [16] Andrea Frome, Gregory S. Corrado, Jonathon Shlens, Samy Bengio, Jeffrey Dean, Marc’Aurelio Ranzato, and Tomas Mikolov. DeViSE: A deep visual-semantic embedding model. In *Proc. NeurIPS*, 2013.
- [17] Weifeng Ge, Xiangru Lin, and Yizhou Yu. Weakly supervised complementary parts models for fine-grained image classification from the bottom up. In *CVPR*, pages 3034–3043, 2019.
- [18] Timnit Gebru, Judy Hoffman, and Li Fei-Fei. Fine-grained recognition in the wild: A multi-task domain adaptation approach. In *ICCV*, pages 1349–1358, 2017.
- [19] Xiangteng He and Yuxin Peng. Fine-grained image classification via combining vision and language. In *CVPR*, pages 5994–6002, 2017.
- [20] Xiangteng He, Yuxin Peng, and Liu Xie. A new benchmark and approach for fine-grained cross-media retrieval. In *ACM Multimedia*, pages 1740–1748, 2019.
- [21] Grant Van Horn, Oisín Mac Aodha, Yang Song, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The iNaturalist challenge 2017 dataset. *arXiv.cs*, abs/1707.06642, 2017.
- [22] Peng Hu, Xu Wang, Liangli Zhen, and Dezhong Peng. Separated variational hashing networks for cross-modal retrieval. In *ACM Multimedia*, pages 1721–1729, 2019.
- [23] Lun Huang, Wenmin Wang, Jie Chen, and Xiao-Yong Wei. Attention on attention for image captioning. In *ICCV*, pages 4634–4643, 2019.
- [24] Zixuan Huang and Yin Li. Interpretable and accurate fine-grained recognition via region grouping. In *CVPR*, pages 8662–8672, 2020.
- [25] Karen Sparck Jones. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 1972.
- [26] Michael Kampffmeyer, Yinbo Chen, Xiaodan Liang, Hao Wang, Yujia Zhang, and Eric P Xing. Rethinking knowledge graph propagation for zero-shot learning. In *CVPR*, pages 11487–11496, 2019.
- [27] Elyor Kodirov, Tao Xiang, and Shaogang Gong. Semantic autoencoder for zero-shot learning. In *CVPR*, pages 3174–3183, 2017.
- [28] Jonathan Krause, Benjamin Sapp, Andrew Howard, Howard Zhou, Alexander Toshev, Tom Duerig, James Philbin, and Li Fei-Fei. The unreasonable effectiveness of noisy data for fine-grained recognition. In *ECCV*, pages 301–320. Springer, 2016.

- [29] Neeraj Kumar, Peter N Belhumeur, Arijit Biswas, David W Jacobs, W John Kress, Ida C Lopez, and João VB Soares. Leafsnap: A computer vision system for automatic plant species identification. In *ECCV*, pages 502–516. Springer, 2012.
- [30] Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. On the sentence embeddings from bert for semantic textual similarity. In *EMNLP*, pages 9119–9130, 2020.
- [31] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [32] Yang Long, Li Liu, Ling Shao, Fumin Shen, Guiguang Ding, and Jungong Han. From zero-shot learning to conventional supervised classification: Unseen visual data synthesis. In *CVPR*, pages 1627–1636, 2017.
- [33] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, pages 13–23, 2019.
- [34] Sean MacAvaney, Andrew Yates, Arman Cohan, and Nazli Goharian. Cedar: Contextualized embeddings for document ranking. In *Proceedings of the 42nd Intl. ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1101–1104, 2019.
- [35] Christopher D Manning, Hinrich Schütze, and Prabhakar Raghavan. *Introduction to information retrieval*. Cambridge university press, 2008.
- [36] Maria-Elena Nilsback and Andrew Zisserman. A visual vocabulary for flower classification. In *CVPR*, volume 2, pages 1447–1454, 2006.
- [37] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics & Image Processing*, pages 722–729. IEEE, 2008.
- [38] Rodrigo Nogueira and Kyunghyun Cho. Passage re-ranking with bert. *arXiv preprint arXiv:1901.04085*, 2019.
- [39] Rodrigo Nogueira, Zhiying Jiang, and Jimmy Lin. Document ranking with a pretrained sequence-to-sequence model. *EMNLP*, 2020.
- [40] Yuxin Peng, Xin Huang, and Jinwei Qi. Cross-media shared representation by hierarchical learning with multiple deep networks. In *IJCAI*, pages 3846–3853, 2016.
- [41] Yuxin Peng, Xin Huang, and Yunzhen Zhao. An overview of cross-media retrieval: Concepts, methodologies, benchmarks, and challenges. *IEEE Transactions on circuits and systems for video technology*, 28(9):2372–2385, 2017.
- [42] Yuxin Peng, Jinwei Qi, Xin Huang, and Yuxin Yuan. Ccl: Cross-modal correlation learning with multigrained fusion by hierarchical network. *IEEE Transactions on Multimedia*, 20(2):405–420, 2017.

- [43] Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Nikhil Rasiwasia, Gert RG Lanckriet, Roger Levy, and Nuno Vasconcelos. On the role of correlation and abstraction in cross-modal multimedia retrieval. *TPAMI*, 36(3):521–535, 2013.
- [44] Ruizhi Qiao, Lingqiao Liu, Chunhua Shen, and Anton Van Den Hengel. Less is more: zero-shot learning from online textual documents with noise suppression. In *CVPR*, pages 2249–2257, 2016.
- [45] 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. *ICML*, 2021.
- [46] Scott Reed, Zeynep Akata, Honglak Lee, and Bernt Schiele. Learning deep representations of fine-grained visual descriptions. In *CVPR*, pages 49–58, 2016.
- [47] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *EMNLP-IJCNLP*, pages 3973–3983, 2019.
- [48] Stephen E Robertson, Steve Walker, Susan Jones, Micheline M Hancock-Beaulieu, Mike Gatford, et al. Okapi at trec-3. *Nist Special Publication Sp*, 109:109, 1995.
- [49] Marcel Simon and Erik Rodner. Neural activation constellations: Unsupervised part model discovery with convolutional networks. In *ICCV*, pages 1143–1151, 2015.
- [50] Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-shot learning through cross-modal transfer. In *NeurIPS*, pages 935–943, 2013.
- [51] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, Marc Proesmans, and Luc Van Gool. Scan: Learning to classify images without labels. In *ECCV*, 2020.
- [52] Grant Van Horn, Steve Branson, Ryan Farrell, Scott Haber, Jessie Barry, Panos Ipeirotis, Pietro Perona, and Serge Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *CVPR*, pages 595–604, 2015.
- [53] 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*, pages 8769–8778, 2018.
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [55] Andrea Vedaldi, Siddharth Mahendran, Stavros Tsogkas, Subhransu Maji, Ross Girshick, Juho Kannala, Esa Rahtu, Iasonas Kokkinos, Matthew B Blaschko, David Weiss, et al. Understanding objects in detail with fine-grained attributes. In *CVPR*, pages 3622–3629, 2014.
- [56] Maunil R. Vyas, Hemanth Venkateswara, and Sethuraman Panchanathan. Leveraging seen and unseen semantic relationships for generative zero-shot learning. [abs/2007.09549](https://arxiv.org/abs/2007.09549), 2020.

- [57] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [58] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. Adversarial cross-modal retrieval. In *ACM Multimedia*, pages 154–162, 2017.
- [59] Jiuniu Wang, Wenjia Xu, Qingzhong Wang, and Antoni B Chan. Compare and reweight: Distinctive image captioning using similar images sets. In *ECCV*, 2020.
- [60] Josiah Wang, Katja Markert, and Mark Everingham. Learning models for object recognition from natural language descriptions. In *BMVC*, 2009.
- [61] Qingzhong Wang, Jia Wan, and Antoni B. Chan. On diversity in image captioning: Metrics and methods. *TPAMI*, PP, 2020.
- [62] Weiran Wang and Karen Livescu. Large-scale approximate kernel canonical correlation analysis. *ICLR*, 2016.
- [63] Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, 2018.
- [64] Yongqin Xian, Zeynep Akata, Gaurav Sharma, Quynh Nguyen, Matthias Hein, and Bernt Schiele. Latent embeddings for zero-shot classification. In *CVPR*, pages 69–77, 2016.
- [65] Yongqin Xian, Christoph H Lampert, Bernt Schiele, and Zeynep Akata. Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *TPAMI*, 41(9): 2251–2265, 2018.
- [66] Yongqin Xian, Tobias Lorenz, Bernt Schiele, and Zeynep Akata. Feature generating networks for zero-shot learning. In *CVPR*, pages 5542–5551, 2018.
- [67] Yongqin Xian, Saurabh Sharma, Bernt Schiele, and Zeynep Akata. f-vaegan-d2: A feature generating framework for any-shot learning. In *CVPR*, pages 10275–10284, 2019.
- [68] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, pages 2048–2057, 2015.
- [69] Zhe Xu, Shaoli Huang, Ya Zhang, and Dacheng Tao. Webly-supervised fine-grained visual categorization via deep domain adaptation. *TPAMI*, 40(5):1100–1113, 2016.
- [70] Zeynep Akkalyoncu Yilmaz, Wei Yang, Haotian Zhang, and Jimmy Lin. Cross-domain modeling of sentence-level evidence for document retrieval. In *EMNLP-IJCNLP*, pages 3481–3487, 2019.
- [71] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based r-cnns for fine-grained category detection. In *ECCV*, pages 834–849. Springer, 2014.

- [72] Liangli Zhen, Peng Hu, Xu Wang, and Dezhong Peng. Deep supervised cross-modal retrieval. In *Proc. CVPR*, pages 10394–10403, 2019.
- [73] Heliang Zheng, Jianlong Fu, Tao Mei, and Jiebo Luo. Learning multi-attention convolutional neural network for fine-grained image recognition. In *ICCV*, pages 5209–5217, 2017.
- [74] Yizhe Zhu, Mohamed Elhoseiny, Bingchen Liu, Xi Peng, and Ahmed Elgammal. A generative adversarial approach for zero-shot learning from noisy texts. In *CVPR*, pages 1004–1013, 2018.