

LTGC: Long-tail Recognition via Leveraging LLMs-driven Generated Content

Qihao Zhao^{1,2,*}, Yalun Dai^{3,*}, Hao Li⁴, Wei Hu¹, Fan Zhang^{1†}, Jun Liu²

¹Beijing University of Chemical Technology, China

²Singapore University of Technology and Design, Singapore

³Nanyang Technological University, Singapore

⁴Northwestern Polytechnical University, China

Abstract

Long-tail recognition is challenging because it requires the model to learn good representations from tail categories and address imbalances across all categories. In this paper, we propose a novel generative and fine-tuning framework, LTGC, to handle long-tail recognition via leveraging generated content. Firstly, inspired by the rich implicit knowledge in large-scale models (e.g., large language models, LLMs), LTGC leverages the power of these models to parse and reason over the original tail data to produce diverse tail-class content. We then propose several novel designs for LTGC to ensure the quality of the generated data and to efficiently fine-tune the model using both the generated and original data. The visualization demonstrates the effectiveness of the generation module in LTGC, which produces accurate and diverse tail data. Additionally, the experimental results demonstrate that our LTGC outperforms existing state-of-the-art methods on popular long-tailed benchmarks.

1. Introduction

In the real world, data often exhibits a long-tailed distribution, posing significant challenges for computer vision recognition [44, 54]. These challenges include (1) Class Imbalance: within the dataset, some classes (termed "head" classes) are abundantly represented, while others (termed "tail" classes) have few samples. This imbalanced distribution during training may cause the model to attend more on the head classes, neglecting the tail classes [42]. (2) Tail Data Scarcity: Data scarcity refers to tail classes having extremely limited samples, which lack diversity and are insufficient to effectively train a model [8, 24]. It prevents a model's ability to learn the feature invariant, which is necessary to recognize these categories correctly [35].

*Equal contribution.

†The Corresponding author is with the College of Information Science and Technology and the Interdisciplinary Research Center for Artificial Intelligence, Beijing University of Chemical Technology, China

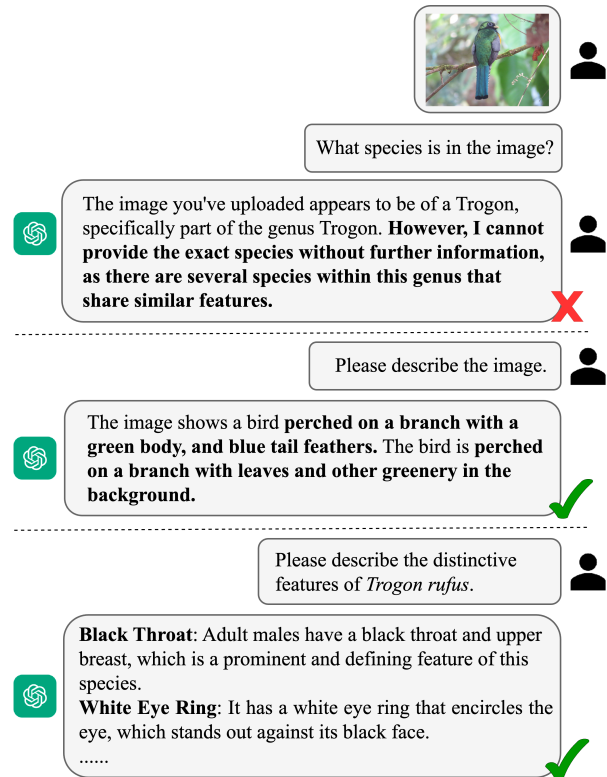


Figure 1. **Top:** Here, our LLMs use ChatGPT. For the *Trogon Rufus* category, when we asked ChatGPT, "What species is in the picture?" we did not get the expected answer due to the complexity of the question. **Middle and Down:** In contrast, when we asked some easy questions for the ChatGPT, "Please describe the image." or "Please describe the distinctive features of *Trogon rufus*." It could accurately answer these questions.

To address these challenges, numerous approaches have emerged within the field of long-tailed recognition, such as re-sampling [5, 19], loss re-weighting [1, 6, 12, 23, 26, 27, 29, 39, 48], ensemble learning [42, 47, 56], decoupling [20], and contrastive learning [11, 25, 40, 59]. Their primary objective is to balance the decision boundaries and parameter weights of the model to improve long-tail recognition.

However, these methods encounter bottlenecks due to the scarcity and limited diversity features of tail-class data [24]. To obtain more diverse representations for tail classes, two types of methods appear: Some methods increase tail-class diversity through data augmentation [9, 24, 52, 55]. Others increase tail-class diversity by transferring features from related classes [8] or large pre-trained models (e.g., CLIP [32]). These transfer learning methods [13, 28, 36] based on CLIP have recently shown great potential in boosting long-tail recognition. The LPT [13] fine-tunes pre-trained models to adapt to target datasets. The VL-LTR [36] leverages text features from CLIP to augment the learning of image features. However, these methods find it hard to obtain the correct and desired knowledge for tail categories. For example, in the real-world data iNaturalist 2018 [17], the category *Aquilegia Pubescens* is characterized only by white or pale yellow colors and spurred shapes. If we transfer similar category features or blindly augment semantic information, such as red color semantics, this category will become more confusing.

Recently, Large Language Models (LLMs) (e.g., ChatGPT) and Large Multimodal Models (LMMs) (e.g., GPT-4V (ChatGPT with Vision [50]) and Minigt4 [7, 58]) due to their wealth of potential knowledge, have been leveraged for a variety of downstream tasks, such as robot task plans [34], open-set object recognition [30], automated robot learning [43] and various visual reasoning tasks [50, 51, 60]. However, due to the bias of imbalanced training data, LMMs perform poorly on some complex tasks [10]. We also find this dilemma in long-tail recognition as shown in Fig. 1. When we asked LMMs about an image in the category of *Trogon rufus*: "What species is in the image?", we did not get the expected answer. In contrast, when we ask some easy questions about this image: "Please describe the image." or "Please describe the distinctive features of *Trogon Rufus*", we get the desired answer. This suggests that although large models underperform in long-tail recognition (e.g., it is difficult to align complex image features with labels.), they still contain desired potential knowledge (e.g., correctly describing images and providing the knowledge of species). Inspired by this, we aim to leverage the abundant potential knowledge of several large models to address the challenges of long-tailed data scarcity and perform long-tailed recognition well. Nevertheless, effectively leveraging large models to encounter these challenges is not trivial: (1) The potential nature of knowledge within large models makes it difficult to extract the desired knowledge that facilitates long-tail recognition tasks. (2) Due to the domain gap between the generated data and the original data [38], effectively using these hybrid data remains an open issue.

To this end, we propose a novel Long-Tail recognition framework via Generated Content, denoted as LTGC,

which is illustrated in Fig. 2. Our LTGC aims to leverage the capabilities of large models for generating explicitly diverse content [15] tailored to the long-tail classes and incorporate novel designs to enhance long-tail recognition. Firstly, motivated by the wealth of potential knowledge in large models and the fact that texts are more controllable, we employ them to expand the tail classes. Specifically, to produce more *diverse and controllable* tail data, we follow the rule of analyzing existing features before generating absent features. To achieve this, LTGC utilizes LMMs to analyze images already in the tail classes to obtain the existing tail-class descriptions list. Then LTGC inputs the existing descriptions list to LLMs to obtain the desired extended tail-class descriptions list for images that are absent from the existing tail classes. Secondly, LTGC utilizes the text-to-image (T2I) model (e.g., DALL-E [33]) to transform these textual descriptions into images. Moreover, inspired by the benefits of chain-of-thought [46], we propose a self-reflection and iterative evaluation module for this process to ensure the diversity and quality of the generated content. Finally, inspired by the advantages of Mixup [52] in merging different image domains, we propose the **BalanceMix** module to address domain shifts of generated images for the fine-tuning process.

In summary, the contributions of our work are as follows:

1. In the **first** time, we propose a novel framework via generated content, LTGC, which leverages the power of large models to address long-tail recognition challenges.
2. We design a series of novel modules to tackle the tail-class image scarcity problem and design the BalanceMix module to efficiently fine-tune the model using the generated data and the original data.
3. Experimental results demonstrate that our LTGC outperforms existing state-of-the-art methods on popular long-tail benchmarks. Also, the visualization illustrates the diversity and controllability of our generated tail images.

2. Related Work

Long-Tail recognition. Many approaches have been proposed in recent years to address the long-tail recognition challenge, such as resampling [5, 19], loss rebalancing [1, 6, 12, 23, 26, 27, 29, 39, 48] to increase the model's focus on the tail categories, ensemble learning [42, 53, 56] to improve the ability to recognize tail samples by enhancing model performance, contrast learning [11, 21, 25, 40, 49, 59] and calibration learning [9, 57] to improve the model's decision boundaries, but due to the scarcity of data on the tails, the model still doesn't have enough features to learn class representations for long-tail recognition. Recently, data augmentation-based methods [8, 9, 24] proposed to handle scarcity-tailed data by augmenting from related features. However, these methods find it difficult to obtain the correct knowledge for tail classes. Other approaches

[13, 36] have achieved good results by transfer learning [37] the feature of the visual-language pre-training model (e.g., CLIP [32]) for long-tailed recognition.

Different from the previous methods, in this paper, from a novel perspective, we explore a generative and fine-tuning framework to handle long-tail recognition by leveraging the rich potential knowledge of large models. Specifically, we design several novel modules to leverage the correct knowledge for generating tail data and perform efficient fine-tuning with the generated data to boost long-tail recognition.

Large Models. Recently, several large models have emerged, including large language models (e.g., ChatGPT [50]), large multi-modal models (e.g., GPT-4 (ision) [50], MiniGPT [58]) and large generated models (e.g., DALL-E [33]). These expansive models are repositories of extensive knowledge [15] and have undergone scrutiny in diverse applications, such as robot task plans [34], open-set object recognition [30], contextual object detection [51], video generation [22], and few-shot segmentation [60]. In this work, we leverage the rich potential knowledge of these large models to handle the long-tail recognition challenge.

3. Method

To perform long-tail recognition well, the challenge is making the model learn diverse representations from tail categories and address imbalances across all categories. In this paper, we propose a novel long-tail recognition framework, LTGC, to handle these challenges. As depicted in Fig. 2, LTGC employ potential knowledge from various off-the-shelf large-scale models to generate and iteratively assess the quality and diversity of tail classes (Sec. 3.1). Furthermore, LTGC proposes the BalanceMix module to facilitating the fine-tuning process with the generated and original datasets (Sec. 3.2). These modules are described in detail below.

3.1. Diverse Tail Images Generation

To learn diversity representations from tail categories, the previous methods have used data augmentation [9, 24] or transfer learning [13, 36]. However, these works find it difficult to obtain the correct and desired diversity knowledge for tail categories. Inspired by the common-sense knowledge in the LLMs and the fact that textual descriptions are more controllable [33, 50], LTGC takes advantage of these to control the detail and diversity of the generated tail classes. Firstly, to better generate diverse images and control the image detail, LTGC aims to generate images that are absent in the original tail data and represent these images in textual form. Specifically, LTGC employs LMMs to analyze the original tail data and obtain the **existing tail-class descriptions list**. Then LTGC leverages the common-sense knowledge of LLMs to **obtain extended tail-class descrip-**

tions based on the existing tail-class descriptions list. Finally, as images are more suitable as training data for long-tail recognition tasks, we utilize the text-to-image module to generate diverse images based on these tail-class descriptions.

3.1.1 Obtaining Existing Tail-class Descriptions List

In order to ensure the diversity of tail-class images, we first analyze the feature information of the original tail-class images before generating new ones. It guarantees that the content of the generated images is distinct from that of the existing tail-class images. Moreover, due to textual descriptions being more controllable [33], we utilize the textual descriptions to control the detail and diversity of the generated images. To achieve this, our LTGC employs Language Model Multimodals (LMMs), such as GPT-4V (Vision) [50], to analyze the feature information and extract textual descriptions of existing tail classes. During this process, the textual responses from LMMs could be varied and sometimes redundant, which may impede the generation of the desired images. Therefore, we employ textual templates to constrain the responses of LMMs, aiming to unify the textual description formats. Inspired by an image of an object that can be fully described or generated by its class and a list of features [4, 41], we design the textual template to include the given class and its features. Furthermore, the introduction of variations in the scenes plays a crucial role in enhancing the model’s ability to generalize [2], a problem that’s particularly acute in long-tailed datasets [35]. Therefore, the textual information of scenes is also important for image generation.

To this end, we design the LMMs’ response **Template 1** for a given class y as follows: "A photo of the class [y], `{with distinctive features}``{in specific scenes}`.". which include the given class, its distinctive features, and specific scenes or environments. With the response template, for a given M_y number class y , we sequentially feed these tail-class images into LMMs along with an instructing **[Prompt 1]** to analyze the features of these images: "Please use the **Template 1** to briefly describe the image of the class [y]." This process is illustrated in Fig. 3, where y is the given class label. As shown, by formulating the description using the template, LMMs would automatically replace the class, features, and background in the response. After performing this process for all images of each class, we compile a list of text descriptions corresponding to each class. Then we intend to extend the tail-class descriptions list with the existing tail-class descriptions list to generate the images lacking in the tail classes.




Figure 2. **Overall framework of LTGC.** LTGC first employs LLMs to analyze the existing tail data to obtain the existing tail-class descriptions list. Then it inputs the list into LLMs to analyze the absent features of the tail classes and employs the T2I model to generate diverse images. Moreover, our designed self-reflection and iterative evaluation modules ensure the diversity and quality of the tail data. Finally, LTGC employs the BalanceMix module to fine-tune the CLIP’s visual encoder with the extended and original data.

3.1.2 Obtaining Extended Tail-class Descriptions List

In this section, we aim to analyze features missing from tail classes and enrich the descriptions of these classes based on the existing tail-class descriptions list. To accomplish this, inspired by the rich common-sense knowledge of Large Language Models (LLMs), such as ChatGPT [50], we extend the tail-class descriptions list through a two-step process: 1) Inputting the existing descriptions list into LLMs, and 2) Designing the **Prompt 2** to guide LLMs in generating the desired descriptions for images that are absent in the given tail class y : "Besides these descriptions mentioned above, please use the **Template 1** to list other possible **{distinctive features}** and **{specific scenes}** for the class $[y]$," which is illustrated in Fig. 4. For each class, we repeat the above two-step process, and then we obtain the tail-class descriptions list for all tail classes.

In addition, in order for these generated descriptions to better complement each tail class, we encourage LLMs to generate descriptions of sufficient number and diversity. To meet this goal, we introduce a **self-reflection** module in this process, aiming to guide LLMs in rethinking if there are any features or scenes that are missed or repeated. It includes two key designs: a number-checking module and a repetition-checking module. (1) The number-checking module aims to guide LLMs in rethinking whether there

are other missing features or scenarios. To achieve this, after posing the initial **[Prompt 2]** for class y , we update the extended descriptions list and re-ask LLMs the **[Prompt 2]** question, incorporating the newly acquired list. For each class y , this iterative process of number-checking continues until a maximum number K_y of the tail class is achieved, where $K_y = M_y + N_y$ and N_y is the number of generated descriptions for class y . (2) The repetition-checking module aims to guide LLMs in rethinking if there are other features or scenes that are repeated at the end of the number-checking iteration. Specifically, we input the extended descriptions list and the following **[Prompt 3]** of each class y for LLMs’ repetition checking: "Please exclude any repetitive **{distinctive features}** and **{specific scenes}** for class $[y]$ in this descriptions list." After that, LLM will filter out the ones with repeated features based on this prompt and return a new list of descriptions. Through the implementation of this two-step process and LLMs’ self-reflection module, we have obtained a reduction in repeated descriptions and an increase in the diversity of the tail-class descriptions for each class.

3.1.3 Transform Descriptions to Images

Above, we obtain the extended tail-class descriptions list for each class y , denoted as L_y . As the images are better adapted to perform visual recognition tasks, in this section,




Figure 3. **Example of the instruction for LLMs.** When both images from tail classes and textual templates are input into LLMs, textual descriptions corresponding to the images can be obtained. By repeatedly performing this operation on the training data, we convert abstract image descriptions into concrete textual descriptions. Finally, we acquire the current textual descriptions list corresponding to each class.

we aim to leverage the image-generative ability of the text-to-image (T2I) method to generate the images from the tail-class descriptions list. In detail, we employ T2I to generate images based on the descriptions list, denoted as:

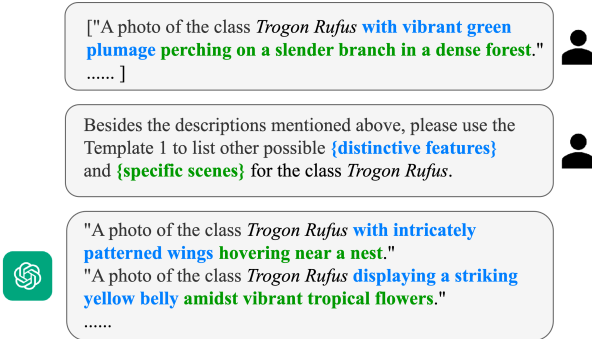


Figure 4. **Example of the instruction for LLMs.** LTGC inputs the existing textual descriptions list to LLMs, which continually extends it with new distinctive features and scene information. During multiple iterations, LTGC generates a new extended textual descriptions list for each class.

$$i_n^y = \text{T2I}(d_n^y), \text{ where } n \in \{1, \dots, N\} \quad (1)$$

where $n \in N$ denotes the n -th generated descriptions for class y , and i_n^y denotes the image generated by T2I model based on d_n^y . However, some generated images may not be of sufficiently high quality to accurately represent their desired classes, as the *lower-quality* images. As shown in Fig. 7, *lower-quality* images will output poorly distinctive features, resulting in more confusion for the class. Therefore, using these generated images directly in training may result in disrupting the model’s prediction of the tail classes.

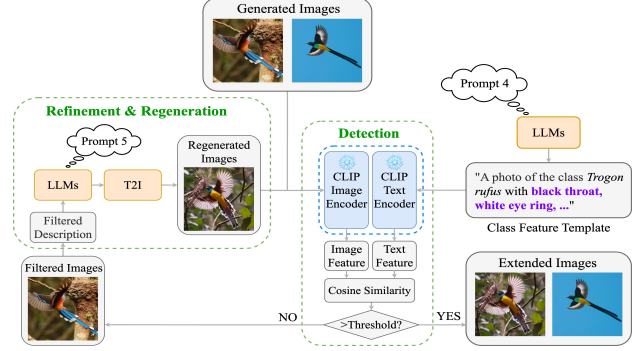


Figure 5. **Illustration of the proposed iterative evaluation module framework.** This module detects lower-quality images through the similarity score \mathcal{S} computed by images and their corresponding class feature template. Then the textual descriptions corresponding to *lower-quality* images are re-input into LLMs for refinement. Finally, the refined textual descriptions are fed into the T2I model for regeneration.

To address the challenge of *lower-quality* images that fail to represent their desired classes accurately, we aim to automatically detect these images, refine their textual descriptions, and regenerate them accordingly. To achieve this, drawing inspiration from the idea that humans improve their understanding through peer feedback, we investigate the possibility of a large model also enhancing its output by integrating feedback from another fundamental model. Therefore, we propose a **iterative evaluation** module, as illustrated in Fig. 5. This module utilizes the fundamental model, CLIP, to serve as the ‘feedback giver’ selected for its proficiency in understanding images and effectively connecting them with textual contexts [32]. It provides feedback about the *lower-quality* images to LLMs, which assists in refining their corresponding descriptions.

Furthermore, to facilitate CLIP in grasping the represented features in images, a straightforward way is to check the degree of match between the image and its class template: “A photo of a [y].” However, due to the fact that CLIP’s text encoder may contain less feature knowledge of the tail classes (e.g., as presented on the iNaturalist 2018 dataset with poor zero-shot recognition performance [13]), we aim to design a **new** textual feature template for each class, highlighting its unique features. Then CLIP employs it to determine the correspondence between the template and the image. Drawing inspiration from the chain-of-thought [30, 46] concept, which suggests that LLMs perform better when provided with additional clues, we further guide LLMs in summarizing the most distinctive features of each class. Specifically, after obtaining the extended tail-class descriptions list, we obtain the **class feature template** C_y that contains the most distinctive features by the following [Prompt 4] and Template 2 for the given class y : [Prompt 4]: “Please use

Template 2 to summarize the most distinctive features of class $[y]$." **Template 2:** "A photo of the class $[y]$ with {feature 1}{feature 2}{...}." This process could guide the LMMs to rethink and summarize the most distinctive features and produce the class feature template C_y for each class y . Furthermore, since we focus on checking whether the feature information of an image could represent its desired class, we do not include scene information in the class feature template. Then we introduce the details of the iterative evaluation module.

As shown in Fig. 5, our iterative evaluation module conducts the following three steps iteratively: **(1) Detection.** To identify *lower-quality* images, we first employ utilize the strong capability of aligning images with text of CLIP to match the feature of generated images i_n^y and its class feature template C_y by cosine similarity metric as follows:

$$S = \text{Encoder}_{\text{vis}}(i_n^y) \cdot \text{Encoder}_{\text{text}}(C_y), \quad (2)$$

where $\text{Encoder}_{\text{vis}}$ denotes the CLIP’s visual encoder, and $\text{Encoder}_{\text{text}}$ denotes the text encoder of the CLIP . The image i_n^y is detected as *lower quality* and filtered out if its similarity score S is below a threshold μ . **(2) Refinement.** Next, also inspired by chain-of-thought [46], we propose the following prompt for refine its descriptions. Specifically, if i_n^y is identified as a *lower quality* image, we prompt LLMs to refine its corresponding description d_n^y to more accurately represent the intended class y , drawing on feedback provided by CLIP, i.e., **[Prompt 5]:** "This description d_n^y doesn't seem to be representative of the class $[y]$. Could you refine it to enhance the distinctive features of class $[y]$?" **(3) Re-generation.** Finally, the image i_n^y is regenerated by the T2I model according to the improved textual description.

By iteratively applying these three stages, we employ CLIP in each cycle to furnish precise feedback to the LLMs, thereby steering the refinement of descriptions to better align with our desired direction. These refined descriptions ensure the production of images that more accurately embody the characteristics of each class.

3.2. BalanceMix

With generated tailed images, the final problem is how to efficiently use these generated data and original data to perform long-tailed recognition well. Due to the domain gap between the generated data and the original long-tail data [38], we propose a method named BalanceMix to handle this challenge. We first define the original data and generated data as \mathcal{D}_o and \mathcal{D}_g . Then BalanceMix balance-sample [5] an image x_i from \mathcal{D}_o and sample an image x_j from \mathcal{D}_g . Meanwhile, it mixes the images x_i and x_j and their corre-

sponding labels, denoted as:

$$\tilde{x} = \lambda \odot x_i + (\mathbf{1} - \lambda) \odot x_j, \quad (3)$$

$$\tilde{y} = \lambda \odot y_i + (\mathbf{1} - \lambda) \odot y_j, \quad (4)$$

where λ in Beta (0,1) distribution. Finally, we fine-tune the CLIP’s vision encoder with LORA [18] on all mixed data pairs (\tilde{x}, \tilde{y}) for efficient long-tail recognition.

4. Experiments

We present the experimental results on three widely used datasets in long-tailed recognition, including ImageNet-LT [27], Places-LT [27], and iNaturalist 2018 [17]. Moreover, we undertake ablation studies specifically on the ImageNet-LT and iNaturalist 2018 datasets to gain deeper insights into the performance of our method. The experimental results of the comparison methods are taken from their original paper, and our results are averaged over three experiments.

4.1. Implementation details.

Evaluation Setup. In all experiments, we evaluate and report top-1 accuracy on their corresponding test set. We also report accuracy on three splits of the classes: Many-shot (more than 100 images), Medium-shot (20 to 100 images), and Few-shot (less than 20 images) [20].

Method Implementation. In our LTGC, we incorporate diverse and specialized knowledge from the off-the-shelf large models. Specifically, for LMM, we use the GPT-4V (ision) [50] version of ChatGPT. For LLM, we use the GPT-4 version of ChatGPT. For T2I, we use DALL-E [33]. For the pre-trained CLIP [32], we use ViT-B/32 [14] for its visual encoder and the architecture introduced in [31] as the text encoder. In LLM’s self-reflection module, we set the maximum number K_y to 100, 300, and 800 for iNaturalist 2018, ImageNet-LT, and Place-LT, respectively. In the iterative evaluation module, the threshold μ is set at 0.8 for ImageNet-LT and Place-LT, and at 0.6 for iNaturalist. In Appendix, we provide a detailed ablation and discussion on the choice of LMM and LLM models, as well as parameters used in the self-reflection process.

4.2. Comparisons with SOTA on Benchmarks

In this section, we compare our proposed LTGC model with state-of-the-art (SOTA) methods on three benchmarks, including Imagenet-LT [27], Places-LT [27], and iNaturalist 2018 [17]. To ensure a fair comparison, we primarily focus on methods based on CLIP, as they also leverage knowledge pretrained on large-scale datasets. These baselines include CLIP Zero-Shot [32, 36] and CLIP Finetune [36], as well as CLIP-based long-tail recognition approaches such as VL-LTR [36], LPT [13], and RAC [28]. Additionally,

we report comparisons with traditional methods (without CLIP) on the challenging and fine-grained, large-scale long-tail dataset, iNaturalist 2018. These have proven the effectiveness of our method in different scenarios.

Results on Imagenet-LT. In Tab. 1, we observe that our LTGC models are superior to other CLIP-based LT methods. For example, our method achieves an overall accuracy of 80.6%, which is 3.4% higher than the existing SOTA method VL-LTR[36]. Moreover, the overall accuracy of our method marginally surpasses the results on the full ImageNet (i.e., 80% [36]).

Results on Places-LT. Tab. 1 shows that compared to other CLIP variant methods, LTGC achieves an overall accuracy of 54.1% and a few-shot accuracy of 52.1%, surpassing the LPT [13] by 4.0% and 5.2%, respectively. Even compared with VL-LTR [36] and RAC [28], which have extra data in training and testing, our LTGC achieves remarkable results.

Results on iNaturalist 2018. In this part, we employ LTGC on a large and fine-grained dataset, iNaturalist 2018. Tab. 2 presents the quantitative results. LTGC leverages the rich knowledge of LLMs and significantly outperforms traditional deep-learning approaches for long-tail recognition. In addition, LTGC attains an overall accuracy of 82.5% and a few-shot accuracy of 82.6%, outperforming all existing SOTA methods based on CLIP. In particular, LTGC also surpasses the retrieval augmented method, RAC [28], by 2.3%.

4.3. Compare with different methods of LMMs.

In the experiments, our LTGC employs GPT-4V (ision) to describe a given image to obtain a text-based feature description. A straightforward way for long-tail recognition is by querying LMMs for the category of the given image. To evaluate the ability of these methods to recognize the long tail images, we constructed the following baselines of LMMs for comparison: MiniGPT4 [58], MiniGPT4-v2 [7], LENS [3], and GPT-4V (ision) [50]. To perform a fair evaluation, we provide the label list [Class 1, Class 2, ..., Class Y] for each dataset, where Y is the number of classes.

The results presented in Tab. 3, show that our method significantly outperforms the baselines that directly query the LMMs. More implementation details are discussed in Appendix.

4.4. Analysis and Ablation Study

Effectiveness of the BalanceMix module. To evaluate the effectiveness of the two proposed designs, we conduct experiments on three different variants: 1) w/o BalanceMix: This variant is directly fine-tuned using the generated data and the original training data. 2) Balanced Sample: This variant utilizes generated data to perform balanced sampling [45] for training. 3) Mixup: This variant employs generated data with Mixup [52] without performing balanced

Table 1. Comparison with SOTA methods on ImageNet-LT and Places-LT.

Dataset	ImageNet-LT		Places-LT	
	Few	All	Few	All
CLIP Zero [36]	58.6	59.8	40.1	38.0
CLIP Finetune [36]	34.5	60.5	22.7	39.7
VL-LTR [36]	59.3	77.2	42.0	50.1
RAC [28]	-	-	41.8	47.2
LPT [13]	-	-	46.9	50.1
LTGC(Ours)	70.5	80.6	52.1	54.1

Table 2. Comparison with SOTA methods on iNaturalist 2018.

Method	Many	Medium	Few	All
Softmax	74.7	66.3	60.0	64.7
LADE [16]	64.4	47.7	34.3	52.3
RIDE [42]	71.5	70.0	71.6	71.8
PaCo [11]	69.5	73.4	73.0	73.0
MDCS [56]	76.5	75.5	75.2	75.6
CLIP Zero [36]	6.1	3.3	2.9	3.4
CLIP Finetune [36]	76.6	74.1	70.2	72.6
VL-LTR [36]	-	-	-	76.8
RAC [28]	75.9	80.5	81.0	80.2
LPT [13]	-	-	79.3	76.1
LTGC(Ours)	77.5	83.9	82.6	82.5

Table 3. Comparison with different LMMs' methods on ImageNet-LT and iNaturalist2018.

Method	ImageNet-LT	iNaturalist 2018
LENS [3]	69.5	17.4
MiniGPT4 [58]	60.4	20.9
MiniGPT4-v2 [7]	68.5	27.1
GPT-4	72.1	64.3
Ours	80.6	82.5

sampling. The results are shown in Tab. 4, demonstrating that our method significantly outperforms the other variants. Although our self-reflection and iterative evaluation modules already ensure the diversity and high quality of the generated images, there still exists a domain gap between the generated and original data [38]. This domain gap could exacerbate the long-tail effect on the test set (discussed in Appendix). A simple balanced sample [45] approach fails to solve the class imbalance caused by the mixing of original and generated data, while the Mixup [52] alone cannot address inter-class imbalances. To this end, our BalanceMix module combines the strengths of both methods, making the generated data well-suited for our framework.

Table 4. Evaluation on the effectiveness of the BalanceMix.

Method	ImageNet-LT	iNaturalist
w/o BalanceMix	58.3	69.5
Balanced sample [45]	63.9	73.8
Mixup [52]	73.4	75.2
Ours	80.6	82.5

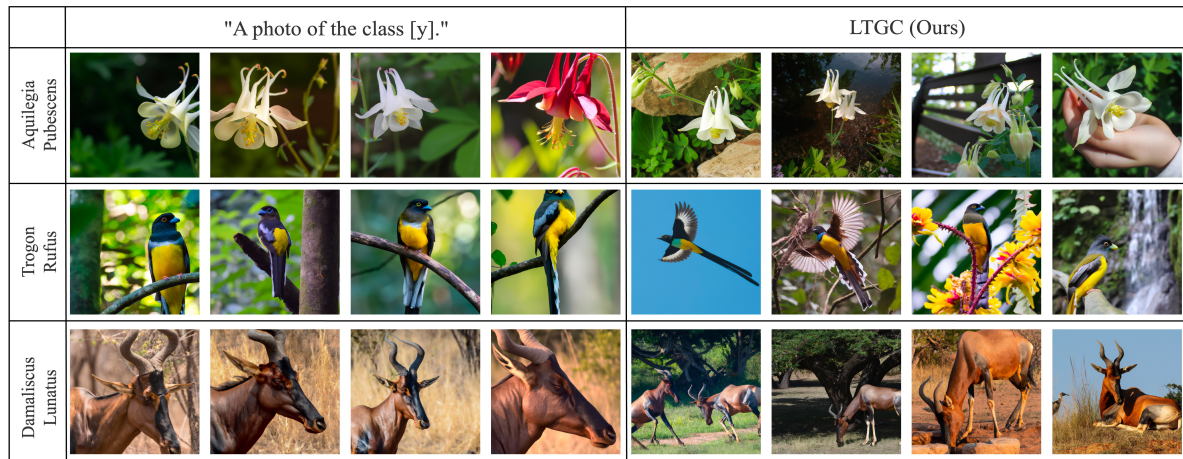


Figure 6. **The visualization of generated images: The template “A photo of the class [y]” and LTGC.** Each row represents a different class. The four images on the left are generated using the simple template ”A photo of the class [y],” which results in images with uniform poses and plain backgrounds. The four images on the right are from the proposed LTGC and demonstrate the diversity of classes.



Figure 7. **The visualization of the images generated before and after passing the iterative evaluation module.** The top row displays images that were filtered out, while the bottom row shows images regenerated by T2I after refining their corresponding descriptions. More visualizations are in the Appendix.

4.5. Visualization

Visualization of Generated Images: Template “A photo of the class [y]” vs our LTGC. Furthermore, we compare the images generated by LTGC with a simple prompt, “A photo of the class [y]” for the T2I model. As Fig. 6 illustrates, LTGC generates more accurate and diverse images compared to those generated by the simple prompt. For example, in the absence of control over the category description, the T2I model generates red features that do not belong to class *Aquilegia Pubescens*. In addition, with rich text for image generation, our approach also generates more diverse and accurate images compared to images generated by a simple template.

Visualization on iterative evaluation module. We compare images before and after refinement using the iterative evaluation module, and the visualization results are shown in Fig. 7. It shows that images before refinement often possess ambiguous semantic information, and the features of the corresponding classes are not distinct. However, after the refinement process, the quality of the images is substantially enhanced, and the distinctive features of the corresponding classes become more pronounced.

5. Conclusion

This paper introduces a novel generative and fine-tuning framework, named LTGC, to address the challenge of long-tail recognition. LTGC leverages the abundant potential knowledge embedded in large-scale models to generate diverse data for tail categories. The framework incorporates innovative designs to ensure the quality of the generated data and to fine-tune the model efficiently using both the generated and original data. The experimental results indicate that LTGC surpasses existing state-of-the-art methods on well-known long-tail benchmarks. In the future, we will explore the robustness of large-scale models for application in other areas, such as long-tail semantic segmentation and dissent detection.

Acknowledgement

This work was supported by the National Natural Science Foundation of China under Grant No.62271034. This project is also supported by the Ministry of Education, Singapore, under the SUTD Kickstarter Initiative Project (SKI 2021.02.06). Additionally, I would like to express my gratitude to Haoxuan Qu and He Huang for their valuable suggestions.

References

- [1] Emanuel Sanchez Aimar, Arvi Jonnarth, Michael Felsberg, and Marco Kuhlmann. Balanced product of calibrated experts for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19967–19977, 2023. 1, 2
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 3
- [3] William Berrios, Gautam Mittal, Tristan Thrush, Douwe Kiela, and Amanpreet Singh. Towards language models that can see: Computer vision through the lens of natural language. *arXiv preprint arXiv:2306.16410*, 2023. 7
- [4] Michel Besserve, Arash Mehrjou, Rémy Sun, and Bernhard Schölkopf. Counterfactuals uncover the modular structure of deep generative models. In *International Conference on Learning Representations*, 2019. 3
- [5] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural Networks*, 106:249–259, 2018. 1, 2, 6
- [6] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *arXiv preprint arXiv:1906.07413*, 2019. 1, 2
- [7] Jun Chen, Deyao Zhu¹ Xiaoqian Shen¹ Xiang Li, Zechun Liu² Pengchuan Zhang, Raghuraman Krishnamoorthi² Vikas Chandra² Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: Large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 2, 7
- [8] Xiaohua Chen, Yucan Zhou, Dayan Wu, Wanqian Zhang, Yu Zhou, Bo Li, and Weiping Wang. Imagine by reasoning: A reasoning-based implicit semantic data augmentation for long-tailed classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 356–364, 2022. 1, 2
- [9] Hsin-Ping Chou, Shih-Chieh Chang, Jia-Yu Pan, Wei Wei, and Da-Cheng Juan. Remix: rebalanced mixup. In *Computer Vision—ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part VI 16*, pages 95–110. Springer, 2020. 2, 3
- [10] Chenhang Cui, Yiyang Zhou, Xinyu Yang, Shirley Wu, Linjun Zhang, James Zou, and Huaxiu Yao. Holistic analysis of hallucination in gpt-4v(ision): Bias and interference challenges, 2023. 2
- [11] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Ji-aya Jia. Parametric contrastive learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 715–724, 2021. 1, 2, 7
- [12] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019. 1, 2
- [13] Bowen Dong, Pan Zhou, Shuicheng Yan, and Wangmeng Zuo. Lpt: Long-tailed prompt tuning for image classification. *arXiv preprint arXiv:2210.01033*, 2022. 2, 3, 5, 6, 7
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 6
- [15] Lin Geng Foo, Hossein Rahmani, and Jun Liu. Ai-generated content (aigc) for various data modalities: A survey, 2023. 2, 3
- [16] Youngkyu Hong, Seungju Han, Kwanghee Choi, Seokjun Seo, Beomsu Kim, and Buru Chang. Disentangling label distribution for long-tailed visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6626–6636, 2021. 7
- [17] Grant Van Horn, Oisín Mac Aodha, Yang Song, Alexander Shepard, Hartwig Adam, Pietro Perona, and Serge J. Belongie. The inaturalist challenge 2017 dataset. *CoRR*, abs/1707.06642, 2017. 2, 6
- [18] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 6
- [19] Nathalie Japkowicz and Shaju Stephen. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5): 429–449, 2002. 1, 2
- [20] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019. 1, 6
- [21] Bingyi Kang, Yu Li, Sa Xie, Zehuan Yuan, and Jiashi Feng. Exploring balanced feature spaces for representation learning. In *International Conference on Learning Representations*, 2020. 2
- [22] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. *arXiv preprint arXiv:2303.13439*, 2023. 3
- [23] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 29(8):3573–3587, 2017. 1, 2
- [24] Shuang Li, Kaixiong Gong, Chi Harold Liu, Yulin Wang, Feng Qiao, and Xinjing Cheng. Metasaug: Meta semantic augmentation for long-tailed visual recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5212–5221, 2021. 1, 2, 3
- [25] Tianhong Li, Peng Cao, Yuan Yuan, Lijie Fan, Yuzhe Yang, Rogerio S Feris, Piotr Indyk, and Dina Katabi. Targeted supervised contrastive learning for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6918–6928, 2022. 1, 2
- [26] Tsung-Yi Lin, Priyal Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. *IEEE*

- transactions on pattern analysis and machine intelligence*, 2018. 1, 2
- [27] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 1, 2, 6
- [28] Alexander Long, Wei Yin, Thalaiyasingam Ajanthan, Vu Nguyen, Pulak Purkait, Ravi Garg, Alan Blair, Chunhua Shen, and Anton van den Hengel. Retrieval augmented classification for long-tail visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6959–6969, 2022. 2, 6, 7
- [29] Aditya Krishna Menon, Sadeep Jayasumana, Ankit Singh Rawat, Himanshu Jain, Andreas Veit, and Sanjiv Kumar. Long-tail learning via logit adjustment. *arXiv preprint arXiv:2007.07314*, 2020. 1, 2
- [30] Haoxuan Qu, Xiaofei Hui, Yujun Cai, and Jun Liu. Lmc: Large model collaboration with cross-assessment for training-free open-set object recognition. *arXiv preprint arXiv:2309.12780*, 2023. 2, 3, 5
- [31] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019. 6
- [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 *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2, 3, 5, 6
- [33] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021. 2, 3, 6
- [34] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. Progprompt: Generating situated robot task plans using large language models. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11523–11530, 2023. 2, 3
- [35] Kaihua Tang, Mingyuan Tao, Jiaxin Qi, Zhenguang Liu, and Hanwang Zhang. Invariant feature learning for generalized long-tailed classification. In *European Conference on Computer Vision*, pages 709–726. Springer, 2022. 1, 3
- [36] Changyao Tian, Wenhai Wang, Xizhou Zhu, Jifeng Dai, and Yu Qiao. VI-ltr: Learning class-wise visual-linguistic representation for long-tailed visual recognition. In *European Conference on Computer Vision*, pages 73–91. Springer, 2022. 2, 3, 6, 7
- [37] Lisa Torrey and Jude Shavlik. Transfer learning. In *Handbook of research on machine learning applications and trends: algorithms, methods, and techniques*, pages 242–264. IGI global, 2010. 3
- [38] Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective data augmentation with diffusion models. *arXiv preprint arXiv:2302.07944*, 2023. 2, 6, 7
- [39] Peng Wang, Kai Han, Xiu-Shen Wei, Lei Zhang, and Lei Wang. Contrastive learning based hybrid networks for long-tailed image classification. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 943–952, 2021. 1, 2
- [40] Peng Wang, Kai Han, Xiu-Shen Wei, Lei Zhang, and Lei Wang. Contrastive learning based hybrid networks for long-tailed image classification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 943–952, 2021. 1, 2
- [41] Tan Wang, Zhongqi Yue, Jianqiang Huang, Qianru Sun, and Hanwang Zhang. Self-supervised learning disentangled group representation as feature. *Advances in Neural Information Processing Systems*, 34:18225–18240, 2021. 3
- [42] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella X Yu. Long-tailed recognition by routing diverse distribution-aware experts. *arXiv preprint arXiv:2010.01809*, 2020. 1, 2, 7
- [43] Yufei Wang, Zhou Xian, Feng Chen, Tsun-Hsuan Wang, Yian Wang, Katerina Fragkiadaki, Zackory Erickson, David Held, and Chuang Gan. Robogen: Towards unleashing infinite data for automated robot learning via generative simulation, 2023. 2
- [44] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 7032–7042, 2017. 1
- [45] Yu-Xiong Wang, Deva Ramanan, and Martial Hebert. Learning to model the tail. *Advances in neural information processing systems*, 30, 2017. 7
- [46] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022. 2, 5, 6
- [47] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*, pages 247–263. Springer, 2020. 1
- [48] Cihang Xie and Alan Yuille. Intriguing properties of adversarial training at scale. *arXiv preprint arXiv:1906.03787*, 2019. 1, 2
- [49] Yuzhe Yang and Zhi Xu. Rethinking the value of labels for improving class-imbalanced learning. *Advances in neural information processing systems*, 33:19290–19301, 2020. 2
- [50] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of llms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 2023. 2, 3, 4, 6, 7
- [51] Yuhang Zang, Wei Li, Jun Han, Kaiyang Zhou, and Chen Change Loy. Contextual object detection with multimodal large language models. *arXiv preprint arXiv:2305.18279*, 2023. 2, 3
- [52] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 2, 7

- [53] Yifan Zhang, Bryan Hooi, Lanqing Hong, and Jiashi Feng. Test-agnostic long-tailed recognition by test-time aggregating diverse experts with self-supervision. *arXiv preprint arXiv:2107.09249*, 2021. [2](#)
- [54] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. [1](#)
- [55] Qihao Zhao, Yangyu Huang, Wei Hu, Fan Zhang, and Jun Liu. Mixpro: Data augmentation with maskmix and progressive attention labeling for vision transformer. In *The Eleventh International Conference on Learning Representations*, 2022. [2](#)
- [56] Qihao Zhao, Chen Jiang, Wei Hu, Fan Zhang, and Jun Liu. Mdcx: More diverse experts with consistency self-distillation for long-tailed recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11597–11608, 2023. [1](#), [2](#), [7](#)
- [57] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16489–16498, 2021. [2](#)
- [58] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [2](#), [3](#), [7](#)
- [59] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2022. [1](#), [2](#)
- [60] Lanyun Zhu, Tianrun Chen, Deyi Ji, Jieping Ye, and Jun Liu. Llafs: When large-language models meet few-shot segmentation, 2023. [2](#), [3](#)